

SPINOFFS AND THE MARKET FOR IDEAS*

BY SATYAJIT CHATTERJEE AND ESTEBAN ROSSI-HANSBERG¹

Federal Reserve Bank of Philadelphia, U.S.A.; Princeton University, U.S.A.

We present a theory of entry through spinoffs where workers generate ideas and possess private information concerning their quality. Because quality is privately observed, adverse selection implies that the market can only offer a price that reflects the average quality of ideas sold. Only workers with good ideas decide to spin off, whereas workers with mediocre ideas sell them. Existing firms pay a price for ideas sold in the market that implies zero expected profits. Hence, firms' project selection is independent of firm size, which can lead to scale-independent growth. This mechanism results in invariant firm-size distributions that resemble the data.

1. INTRODUCTION

The generation and implementation of new ideas shapes industry dynamics and the structure of firms. Ideas can be generated in many different contexts, but many important innovations have been developed by workers of established firms. In some cases, workers sell their ideas to established firms (including their own), and in other cases, they use them to start new firms: spinoffs. Whether an innovation by a worker is implemented by existing firms, leads to a spinoff, or is discarded depends on the initial knowledge about the idea, as well as the profits that the different entities can generate by implementing it. In this article, we present a theory of entry through spinoffs in which the key ingredient is the originator's private information concerning the quality of his new idea. Because quality is privately observed, by the standard adverse-selection logic, the market can at best offer a price that reflects the average quality of ideas sold. This gives the holders of above-average-quality ideas the incentive to spin off. But spinning off is costly and there is a critical quality level at which the loss from selling the idea at the market price balances the costs imposed by spinning off. Ideas that are of higher quality than this critical level lead to spinoffs, whereas ideas that are of lower quality than the critical level are sold to existing firms.

Our theory of spinoffs implies a new model of firm entry and firm growth. Ideas that are spun off generate entry of new businesses, whereas ideas that are sold to existing firms generate new employment in existing firms. We take the view that when a person sets up his or her own business, the person has some new business idea in mind—it could be something as simple as a pizza store in a new location or something sophisticated like new business software. Our model of spinoffs encompasses both cases because the pizza store owner could have “sold” his idea of a new location to a national franchise by partnering with it, and the owner of the new business software could have sold his invention to an existing software company. Thus, in a broad sense, entry occurs when these options to sell new ideas are not exercised, and growth of existing firms occurs when they are exercised. This view also explains our terminology: We call our model a

*Manuscript received October 2009; revised September 2010.

¹ We thank Hal Cole, Ken Hendricks, Hugo Hopenhayn, Boyan Jovanovic, Chris Phelan, Victor Rios-Rull, and participants at several seminars and conferences for useful comments. Rossi-Hansberg thanks the Sloan Foundation for financial support. The views expressed in this article are those of the authors and do not necessarily reflect the views of the Federal Reserve System or the Federal Reserve Bank of Philadelphia. This article is available free of charge at www.philadelphiafed.org/research-and-data/publications/working-papers/. Please address correspondence to: Satyajit Chatterjee, Research Department, Federal Reserve Bank of Philadelphia, Ten Independence Mall, Philadelphia, PA 19106-1574. Phone: 215 574 3861. Fax: 215 574 4303. E-mail: chatterjee.satyajit@gmail.com.

theory of spinoffs precisely because the alternative is to sell the new idea to an established firm. We explore this model of entry and firm growth at some length in the article.

There are several key elements in our theory. The first key element, of course, is private information. Specifically, we assume that a worker with a new idea has private information on the mean payoff from the idea. The worker can decide to either create a new firm to implement the idea or sell the idea to an established firm at a price that is independent of the (privately observed) mean return. If he does the latter, he can credibly reveal the mean return to the buying firm *after* the sale, since there is nothing at stake for him at that point. Knowing the mean return, the buying firm then decides whether to implement the idea. Low quality ideas are discarded without being implemented. Implementing an idea means producing with it for one period. Producing with an idea—either in a spinoff or in an established firm—reveals the actual payoff from the idea. We call an idea that has been implemented, and therefore that has a known payoff, a project. Production requires one unit of labor. Since labor is costly, the project payoff can be low enough to make further use nonoptimal. In that case, the project is dropped. If the project payoff is sufficiently high, the project is run forever and provides a constant source of profits to the entrepreneur who implemented it.

The second key element is the costs of spinning off. We assume that managing one or more projects is a full-time job that leaves the entrepreneur no time to devote to inventing new ideas. Thus, becoming an entrepreneur implies giving up on the possibility of spinning off in the future with an even better idea. Clearly this option to spin off in the future has value and, consequently, prospective entrepreneurs—workers who decide to test an idea on their own—are more choosy about which projects to accept and run than established entrepreneurs. Because project returns are specific to the entrepreneur who tests the idea, projects that are discarded by a prospective entrepreneur cannot be sold to established entrepreneurs. This specificity makes workers with ideas more selective in choosing which ideas to test on their own compared to established entrepreneurs. Thus the mean return at which established entrepreneurs are just willing to test an idea is lower than the mean return that just induces a worker to spin off and test on his own. The gap in these thresholds implies that ideas can sell at a positive price in the market. To use an analogy, it is like having a “lemons problem” in which withholding a used car from the market is costly (because, say, the person has to pay insurance costs, garage fees, etc.), and, therefore, the person may sell a car even if the price fetched in the market is lower than the (privately observed) value of the car.²

The third key element is competition in the market for ideas. When individuals are risk neutral or have constant absolute risk aversion and competition forces all the (expected) surplus from new ideas to go to the seller of the idea, the mean return for which established entrepreneurs are indifferent between testing an idea or discarding it is independent of the size of the entrepreneur’s firm. Hence, in equilibrium, heterogeneous firms are indifferent about how many ideas to buy. If they find out about ideas at a rate proportional to their size, this implies scale-independent growth for established firms, for which there is some evidence in the U.S. data. If the market for ideas was not competitive, firms would not be indifferent about the number of ideas they buy, and scale-independent growth would not be an equilibrium outcome of the model. We show that this growth process, together with the entry process predicted by the model, leads to a realistic size distribution of firms.

There are two strands of literature directly related to our article: the one on spinoffs and the one on firm/industry dynamics. A key contribution of this article is to combine these two,

² It seems plausible that setting up and running a new firm will leave the entrepreneur with little time to innovate, at least for some length of time. For simplicity we go to the extreme and assume that this length of time is infinite. What is fundamental is that there be *some* cost to spinning off. In Section B of the Appendix B, we consider an extension of our model in which project returns are not specific to the entrepreneur who tests the project and entrepreneurs do not forgo the opportunity of getting ideas. However, (new) entrepreneurs must pay a resource cost to start up a new firm. The main results hold in this environment as well.

otherwise separate, literatures.³ Turning first to the literature on spinoffs, our article is related to Anton and Yao (1994, 1995). They study the problem of a worker who privately learns of an innovation and must decide between revealing the innovation to his employer in return for compensation or keeping it secret and exploiting the innovation independently in a spinoff. In their model, the spinoff competes directly with the parent firm and it is the threat of competition that makes it (potentially) rational for the parent firm to compensate the inventor after learning about the innovation. If competition reduces profits enough, the inventor has a credible threat and the equilibrium outcome is for him to reveal the idea in return for adequate compensation. Franco and Filson (2006) provide a theory of spinoffs based on imitation. Established firms understand that workers acquire know-how on the job and eventually become knowledgeable enough to profitably set up a competing firm. In equilibrium, the workers “pay” for this valuable know-how by accepting a lower wage. In contrast to these studies, in our model spinoffs do not compete with established firms, so the threat of competition, or imitation, is absent. Instead, as noted above, the reason ideas get sold at all is that spinning off is costly. This allows us to broaden the scope of the analysis beyond a narrowly defined industry. Silveira and Wright (2010) study the market for ideas in the context of a search model and focus on the role of liquidity provision in the functioning of this market. Neither of the latter two papers focuses on the friction created by private information.

Turning to the literature on firm dynamics, previous studies have mostly taken a different approach. The seminal works of Jovanovic (1982), Hopenhayn (1992), and Ericson and Pakes (1995) study firm dynamics that result from a stream of productivity levels drawn from exogenously specified distributions for existing and new firms. So do more recent papers such as Luttmer (2007), Klette and Kortum (2004), and Rossi-Hansberg and Wright (2007). In contrast to these studies, we stress the fact that new ideas occur to *people* as opposed to organizations. And the people to whom these ideas occur *choose* the organization that gets to implement their ideas—established firms or their own start-ups. In our model, a key implication of this choice is that the distribution from which *established firms* draw their project payoffs (or, equivalently, their productivity shocks) and the distribution from which *new firms* draw their project payoffs are both *endogenously* determined. Furthermore, the distribution of payoffs for new firms has a higher mean than the distribution for established firms. As we discuss later in the article, there is empirical evidence to support this implication.

Finally, our theory connects to an older empirical literature on the firm-size distribution. In an early study, Simon and Bonini (1958) established that the distribution of firm sizes in the United States was well approximated by a Yule distribution. The Yule distribution is a one-parameter distribution that results when new entrants are always of some given size and established firms grow, on average, at a rate that is independent of their size (Gibrat’s Law). As noted earlier, our model is consistent with scale independent growth of established firms and, by assumption, spinoffs start off with one employee (or a constant team size). Thus, our framework can deliver a Yule distribution for the size of business firms. Importantly, our theory provides a microfoundation for the single parameter that governs the shape of the Yule distribution by providing a theory of entry through spinoffs. As we show later in the article, there is evidence that this parameter varies across time and across industries, underscoring the need for such a microfoundation. Of course, the actual size distribution of firms in any economy is the result of a variety of forces, including ideosyncratic productivity shocks, financial constraints, and industry-specific human and physical capital. We abstract from these other forces in order to stress the fact that the idea-selection mechanism we propose leads to an entry process that is consistent with the observed economy-wide firm size distribution and with the variation in these distributions across industries.

³ In an interesting contribution, Anton and Yao (2002) model a market for ideas based upon credible partial disclosure via bond posting. However, they do not model a competitive market for ideas nor do they examine the implications for firm dynamics.

In sum, our article provides an equilibrium theory of firm entry and growth based on private information. Given the fundamental role that business ideas play in the formation and growth of firms, and that the quality of these ideas is usually hard to convey, the lack of private information in theories of firm dynamics was an important gap. This article takes a first step toward filling this gap.

The rest of the article is organized as follows. Section 2 describes the model and establishes some basic results. Section 3 characterizes the selection of ideas and projects and the price of ideas when individuals are risk neutral. Section 4 is devoted to exploring some of the implications of our theory of entry through spinoffs and how those implications stack up against available empirical evidence. Section 5 derives the invariant distribution of firm sizes and compares it to the data on firm sizes for the United States as a whole and for a set of two-digit NAICS industries. Section 6 concludes. In Section A of the Appendix we collect all proofs not included in the main text and show that the main points of our theory carry over to the case where individuals have constant absolute risk aversion. In Section B of the Appendix we extend the model to allow for entrepreneurs to generate ideas and a sunk resource cost of new entry and show that the main results of the article continue to hold.

2. THE MODEL

Agents order consumption according to the following utility function:

$$U(\{c_t\}) = \sum_{t=0}^{\infty} \beta^t u(c_t),$$

where $u(c_t) : \mathbb{R} \rightarrow \mathbb{R}$ is strictly increasing, concave, twice continuously differentiable, and bounded. The boundedness assumption is useful to invoke standard contraction mapping arguments but it is not necessary. In Section 3 we study two particular cases, namely, $u(c) = c$ and $u(c) = -ae^{-bc}$, that do not satisfy the boundedness assumption but for which we can solve the Bellman equation in closed form.

Individuals work in two occupations. They can be entrepreneurs or workers. A worker can work for an entrepreneur at a wage $w > 0$ each period and we assume that there is a perfectly elastic supply of workers at this wage. For simplicity, we abstract from a savings decision so individuals consume what they earn each period. Entrepreneurs earn profits from the projects they own and run, and workers earn wages plus any compensation they receive for ideas they sell to entrepreneurs. As noted already, for the main model we assume that entrepreneurs do not generate ideas, only workers do. A worker can become an entrepreneur if he has an idea and decides to spin off and start his own firm.

An idea is a nonreplicable technology to produce consumption goods using labor, specifically, an idea uses one unit of labor.⁴ Consider an entrepreneur who owns a firm with $N \in \{1, 2, \dots\}$ implemented ideas or projects. Then his one-period profits are given by

$$\pi(S, N) = N(S - w),$$

where $S = \frac{1}{N} \sum_{i=1}^N P_i$ denotes the average revenue and P_i is the per period income generated from a particular idea. We assume that $P_i > 0$ with probability one.

We assume that in each period the probability of a worker getting an idea is λ . An entrepreneur does not get ideas but buys ideas from workers. An entrepreneur learns about ideas with probability $0 < \gamma(\lambda, N) < 1$. For now, we do not take a stand on the specification of $\gamma(\cdot)$ but assume, as seems natural, that the probability of learning about an idea in any given period is

⁴ We assume that ideas are nonreplicable technologies in order to determine the scale of each project. If technologies are replicable, we would need a demand structure and goods differentiation to limit the size of each project. This simple extension would complicate our framework without providing new insights.

increasing in λ and N . We will have more to say about the specification of $\gamma(\cdot)$ at the beginning of Section 4. In particular, in equilibrium, the number of ideas an entrepreneur learns about—the demand for ideas—must be equal to the number of ideas generated by workers, the supply of ideas. As we show below, equilibrium in the market for ideas will determine the average probability of an entrepreneur learning about an idea, but not the distribution of probabilities among them. That is, we show that in equilibrium heterogeneous firms are indifferent about how many ideas to buy. Thus, if we write $\gamma(\lambda, N) = \theta\tilde{\gamma}(\lambda, N)$, θ will be an equilibrium object (the average probability of an entrepreneur learning about an idea) and $\tilde{\gamma}(\cdot)$ (the distribution of probabilities among entrepreneurs) is a primitive of our economy. In Section 4 we specify it as λN .

The mean payoff per period from the idea is μ , which is private information to the originator of the idea.⁵ The mean payoff is drawn from a continuous distribution $H(\mu)$ with $H'(\mu) > 0$ for all $\mu \geq 0$. The actual payoff is drawn from a distribution $F_\mu(P)$, where $\int P dF_\mu(P) = \mu$. The realization of P for a given idea can be discovered by implementing the idea for one period.⁶ We impose the following assumptions on this distribution. First, $\int f(P) dF_\mu(P)$ is increasing in μ for all increasing functions f , $F_\mu(0) = 0$ all μ . Second, F_μ is continuous with respect to μ and $\lim_{\mu \rightarrow \infty} F_\mu(w) = 0$.⁷

As mentioned above, entrepreneurs do not get ideas, as they are involved in the management of their firm, but they can buy ideas from workers.⁸ An entrepreneur who has bought an idea can pay w to try it out for one period and observe the realization of P . If he does, he will use the idea to produce as long as his future expected utility from doing so is greater than from dropping it. Entrepreneurs may decide to implement a project even if the stream of profits is negative ($P < w$), since having an extra project may alter the number of ideas they learn about in the future.

A worker who has had an idea this period has two potential uses for it. He can sell his idea to an entrepreneur, in which case he reveals the mean payoff to the entrepreneur who buys it. In this case he earns a wage w plus the price Z at which he sells the idea. The idea is then owned by the entrepreneur and he decides to try it out or not. The worker can also leave with the idea and become an entrepreneur of a firm with only this idea: a spinoff. Note that in the market of ideas, the price of an idea has to be noncontingent on the quality of the idea. The reason is that any contingent contract would give the worker an incentive to lie about the quality of the idea. So the only incentive-compatible price is independent of quality, in which case the agent is indifferent between revealing the true quality of the idea or not. Since this information is useful for the entrepreneur, we assume that the worker does reveal the true quality. The price

⁵ We assume anyone can tell apart genuine ideas from fake ideas. In other words, it is verifiable whether an idea will pay a strictly positive revenue stream with probability one. Without this assumption we could not impose conditions on the realization of the payoffs from these ideas that are common knowledge to workers and entrepreneurs. We can also invoke a different argument. We model the generation of ideas through the parameter λ , which governs the frequency with which workers have ideas. We view this as a shortcut for a model in which it is costly to generate ideas, even fake ones (which pay zero). If the cost of generating fake ideas is not too low, a market for ideas will still form, since the arguments in the article are developed for general λ and distributions H and F_μ .

⁶ We assume two layers of uncertainty (about μ and P) in order to avoid contracts in which the private information is fully revealed at the contracting stage in exchange for a fee, by threatening dire consequences for misrepresentation. For instance, one could write a contract that says that lying about P is a criminal offense. Then, P would be credibly revealed. In our setup, these contracts are not used, since the inventor does not know the realization P , only its mean μ .

⁷ A sufficient condition for $\int f(P) dF_\mu(P)$ increasing in μ is that a higher μ implies a distribution that first-order stochastic dominates a distribution with a lower μ . Since μ is also the mean of the distribution, this assumption need not be satisfied for all probability distributions. However, it is satisfied by the uniform and the normal distributions (our numerical example uses a uniform distribution). The assumption is needed in order to deliver a well-defined ordering of ideas and is necessary for our selection results reported later. The assumption $\lim_{\mu \rightarrow \infty} F_\mu(w) = 0$ means that, as the mean of the distribution increases, the distribution puts less and less weight on outcomes less than or equal to w . So for high enough μ it is impossible to lose by trying out the idea. Again, this is naturally true for the uniform and normal distributions. This assumption is needed to guarantee that some ideas will always lead to spinoffs.

⁸ We relax this assumption in Section B of the Appendix, where we impose an exogenous resource cost to start a new firm.

of an idea Z is determined in equilibrium, where all entrepreneurs will be indifferent between buying ideas or not.

We assume that the implementation of an idea and the return that it generates are specific to the entrepreneur who tests the idea. Hence, projects are entrepreneur specific (but not worker specific). The notion that some entrepreneur-specific knowledge is used to generate output from the particular implementation of an idea seems plausible. We could imagine a more elaborate setting where a worker who tests an idea can sell his realized project to an established entrepreneur at some loss. In this case, established entrepreneurs would be able to expand by buying ideas and testing them and by buying projects directly from workers who test them but wish to delay becoming entrepreneurs. For simplicity, we chose to make the assumption that without the entrepreneur who implemented the idea, the project has zero value.

We also assume that contracts contingent on the realizations of the project payoff are not possible. The implicit assumption is that contingent contracts come hand in hand with additional agency problems (private information, imperfect enforcement, etc.) that make the use of contingent contracts suboptimal.⁹ We also abstract from financing issues. The friction that leads to a spinoff—the private observability of the mean return—will also make it difficult for an entrepreneur to obtain financing for the project. One possibility is to imagine that these financing hurdles impose additional costs on the spinoff. We show in Section B of the Appendix that including such a resource cost does not change the main results. Thus, we simply abstract from these issues in the main body of the article.

2.1. An Entrepreneur's Problem. Consider the problem of an entrepreneur with average revenue S , coming from N existing projects, who owns one new idea with mean payoff μ . If the entrepreneur tests the idea, his value function is

$$V(\mu, S, N) = \int [u(\pi(S, N) + w + P - Z - w)] dF_\mu(P) \\ + \beta \int \max \left[W \left(\frac{NS + P}{N + 1}, N + 1 \right), W(S, N) \right] dF_\mu(P).$$

This period, his expected utility is the result of consuming the profits from the accumulated used projects $\pi(S, N)$, his wage w , the price paid for the idea Z , and the random realization of profits from the new project $P - w$. Note that the distribution from which P is drawn has expected value μ . Denote by $W(S, N)$ the continuation value of an entrepreneur with N projects with average revenue S . If the entrepreneur uses the project, next period he will manage a firm with $N + 1$ projects and average revenue $(NS + P)/(N + 1)$. If he does not use it, next period his continuation value stays constant at $W(S, N)$. The continuation value (or the value without any new idea) of an entrepreneur with N projects with average revenue S is given by

$$W(S, N) = \gamma(\lambda, N) \int^{\mu_H} \max [V(\mu, S, N), u(\pi(S, N) - Z + w) + \beta W(S, N)] dH(\mu) \\ + (1 - \gamma(\lambda, N) H(\mu_H)) [u(\pi(S, N) + w) + \beta W(S, N)]$$

or

$$W(S, N) = \gamma(\lambda, N) \int^{\mu_H} \max [V(\mu, S, N) - u(\pi(S, N) - Z + w) - \beta W(S, N), 0] dH(\mu) \\ + (1 - \gamma(\lambda, N) H(\mu_H)) [u(\pi(S, N) + w) + \beta W(S, N)] \\ + \gamma(\lambda, N) H(\mu_H) [u(\pi(S, N) - Z + w) + \beta W(S, N)],$$

⁹ Anton and Yao (1995) consider this possibility and show that contingent contracting does not replace the “sale” of ideas or start-ups provided the inventor has limited wealth (or liability).

where μ_H denotes the mean revenue value at which workers leave the firm with their idea.¹⁰ The probability of finding out about an idea next period is $\gamma(\lambda, N)$. Any worker with an idea can leave and set up his own firm. He will do so as long as the idea is good enough, that is $\mu \geq \mu_H$. Hence, ideas get implemented in existing firms only if $\mu < \mu_H$. Given that an idea of expected revenue μ is generated, the value of implementing it is, as discussed above, given by $V(\mu, S, N)$. The value of not implementing the idea is given by $u(\pi(S, N) - Z + w) + \beta W(S, N)$, namely, the utility of consuming profits and wage today and paying the price for the idea plus the same continuation value tomorrow. An idea will not be implemented if it provides a very low expected value. If the entrepreneur does not find out about an idea, or if the idea is good enough to generate a spinoff, the value of the entrepreneur is given by $u(\pi(S, N) + w) + \beta W(S, N)$, since he does not pay for the idea. One of these scenarios happens with probability $1 - \gamma(\lambda, N)H(\mu_H)$.

The next lemma shows that the continuation value $W(S, N)$ exists and is increasing and continuous in average revenue S . We then show in Lemma 2 that the value of an entrepreneur with an idea μ , $V(\mu, S, N)$, is increasing and continuous in the expected value of the idea μ and in the average return S . All proofs are relegated to Section A of the Appendix.

LEMMA 1. $W(S, N)$ exists and is strictly increasing in S .

LEMMA 2. $V(\mu, S, N)$ exists and is strictly increasing and continuous in μ and S .

An entrepreneur will implement an idea with expected revenue μ if

$$V(\mu, S, N) > u(\pi(S, N) - Z + w) + \beta W(S, N).$$

Let $\mu_L(S, N)$ be the value of μ that solves

$$(1) \quad V(\mu_L, S, N) = u(\pi(S, N) - Z + w) + \beta W(S, N).$$

Then an entrepreneur will implement an idea as long as $\mu > \mu_L(S, N)$. Thus, we can rewrite $W(S, N)$ as

$$\begin{aligned} W(S, N) &= \gamma(\lambda, N) \int_{\mu_L(S, N)}^{\mu_H} [V(\mu, S, N) - u(\pi(S, N) - Z + w) - \beta W(S, N)] dH(\mu) \\ &\quad + \gamma(\lambda, N) H(\mu_H) [u(\pi(S, N) - Z + w) + \beta W(S, N)] \\ &\quad + (1 - \gamma(\lambda, N) H(\mu_H)) [u(\pi(S, N) + w) + \beta W(S, N)]. \end{aligned}$$

The next lemma shows that there exists a unique function $\mu_L(S, N)$ that satisfies Equation (1).

LEMMA 3. *There exists a unique function $\mu_L(S, N)$ that satisfies Equation (1).*

2.2. *A Worker's Problem.* The expected utility of a worker with an idea μ who decides to spin off is given by

$$V_0(\mu) = \int u(P) dF_\mu(P) + \beta \int \max[W(P, 1), W_0] dF_\mu(P).$$

¹⁰ Note that we are already assuming that workers spin off when they get an idea with $\mu > \mu_H$. We prove below that this is, in fact, the case. In the meantime, all our arguments remain unaffected if we were to define a set M_H that includes the μ 's for which agents spin off. Then the integrals above would integrate over all values of μ that are not in M_H .

The continuation value of a worker currently working in a firm, W_0 , is then given by

$$W_0 = \lambda \int \max[V_0(\mu), u(w + Z) + \beta W_0] dH(\mu) + (1 - \lambda)[u(w) + \beta W_0].$$

Using arguments similar to the ones used above for V , we can show that $V_0(\mu)$ is strictly increasing in μ . A worker with an idea μ will leave the firm and become an entrepreneur if

$$V_0(\mu) > u(w + Z) + \beta W_0.$$

Let μ_H be the value of μ that solves

$$(2) \quad V_0(\mu_H) = u(w + Z) + \beta W_0.$$

Thus, if $\mu > \mu_H$ the worker will leave his employer and set up a new firm. The continuation value of a worker can therefore be written as

$$W_0 = \lambda \int_{\mu_H} V_0(\mu) dH(\mu) + (1 - \lambda)[u(w) + \beta W_0] + \lambda H(\mu_H)(u(w + Z) + \beta W_0).$$

We show formally below that there exists a unique threshold μ_H . Note also that μ_H is constant and so it is independent of the characteristics of the firm (S, N) in which the agent works.

LEMMA 4. *There exists a unique value μ_H that satisfies Equation (2). Furthermore $V_0(\mu)$ is increasing and continuous in μ .*

We still need to define the realized return needed in order to continue with a project once its return is realized. Define $P_L(N, S)$ as

$$W\left(\frac{NS + P_L(N, S)}{N + 1}, N + 1\right) = W(S, N)$$

and P_H by

$$W(P_H, 1) = W_0.$$

Then a firm keeps the project if the realized return is $P \geq P_L(N, S)$ and a spinoff stays in operation if the realized return on the idea that generated the spinoff is such that $P \geq P_H$. Note that if $P < P_H$, the spinoff will exit and the would-be entrepreneur will return to the labor force as a worker. In that case, given that the implementation of his project was specific to him, the project has zero resale value.

2.3. Equilibrium. A long-run equilibrium of this economy is a distribution of firm sizes δ_N , a list of four thresholds, $\mu_L(\cdot)$, μ_H , $P_L(\cdot)$, and P_H , a price of ideas, Z , and the average probability with which an entrepreneur buys an idea, θ (where θ is given by $\gamma(\lambda, N) = \theta \tilde{\gamma}(\lambda, N)$) such that entrepreneurs solve the problem in Section 2.1, workers solve the problem in Section 2.2 and the price, Z , and the average probability of buying an idea, θ , clears the market for ideas:

$$(3) \quad \lambda \sum_{N=1}^{\infty} (N - 1) \delta_N = \theta \sum_{N=1}^{\infty} \tilde{\gamma}(\lambda, N) \delta_N,$$

where the l.h.s. is the supply of ideas and the r.h.s. is the demand for ideas. As we show in the next section, the price Z will be such that entrepreneurs will be indifferent about how many

ideas to buy. Therefore, market clearing will just require that the number of ideas bought by entrepreneurs be equal to the number of ideas generated by workers. That is, market clearing simply determines the value of θ , but it leaves indeterminate the function $\tilde{\gamma}(\lambda, N)$. Hence, if we want to determine the number of ideas bought by a firm, we need to specify $\tilde{\gamma}(\lambda, N)$ as a primitive of the model. The flexibility to specify $\tilde{\gamma}(\lambda, N)$ is then the direct result of a competitive market for ideas and the resulting equilibrium price Z .

3. CHARACTERIZATION

In this section, we characterize the thresholds on the expected revenue from an idea that determine if an idea is thrown away, implemented by a particular firm, or results in a spinoff. For this, we first assume that the utility function is of the form

$$U(\{c_t\}) = \sum_{t=0}^{\infty} \beta^t u(c_t) = \sum_{t=0}^{\infty} \beta^t c_t.$$

We show in Section A of the Appendix that our main results hold under a CARA utility function as well. The main reason to choose these two utility functions is that we can solve the value of an existing firm in closed form given the additive separability or log additive separability of these utility functions.¹¹

Under the assumption that the utility function is linear we can fully solve this problem in closed form. The first result shows that the threshold $\mu_L(S, N)$ is independent of S and N . $\mu_L(S, N)$ independent of S is implied by risk neutrality (or, in the CARA case below, by the fact that risk does not depend on the level of wealth). $\mu_L(S, N)$ constant in N is the result of the market for ideas. Since workers will sell their ideas to whomever is willing to pay more for them, and there is a relative scarcity of ideas, workers extract all the surplus of an idea, and we can solve for the price of an idea in equilibrium. The proposition also yields the result that in equilibrium $P_L = w$, and so entrepreneurs use all projects that give positive returns. In contrast, $P_H > w$ and so spinoffs use projects that give strictly positive returns. The reason is that new entrepreneurs that start a firm with a project with a low realized return have the option of going back to work for a firm and start a new firm in the future with a better project. The proposition also shows that the threshold for implementing ideas through spinoffs is greater than the one for implementing ideas within the firm, $\mu_L < \mu_H$. This is essentially the result of a positive option value of spinning off in the future and our assumption that project returns are entrepreneur specific. Thus, inventors are more selective with the ideas they implement when they spin off than established firms. This also implies that some ideas do not result in spinoffs and so some firms grow. An industry's growth is then the result of entry through spinoffs and growth in the intensive margin.

PROPOSITION 1. *If $u(c_t) = c_t$, then in equilibrium*

- *the value of the firm is given by $W(S, N) = (\pi(S, N) + w)/(1 - \beta)$,*
- *the value of a worker is given by $W_0 = (w + f_0)/(1 - \beta)$ for some positive constant f_0 ,*
- *$\mu_L(S, N)$ is independent of S and N , and $\mu_L(S, N) < w$,*
- *the thresholds for using a project are given by $P_L(S, N) = w$ and $P_H = w + f_0 > w$,*
- *$\mu_L < \mu_H$, so some ideas are implemented within existing firms and some through spinoffs,*

¹¹ These specifications do not satisfy the boundedness assumption we made in Section 2. However, since in these two cases we can solve the functional equation for $W(S, N)$ analytically, it follows from Theorem 9.12 in Stokey et al. (1989) that the solution is, in fact, optimal.




FIGURE 1

SELECTION OF IDEAS

- and the market price of ideas is given by

$$(4) \quad Z = \frac{1}{H(\mu_H)} \int_{\mu_L}^{\mu_H} \left[\mu - w + \frac{\beta}{1 - \beta} \int \max [P - w, 0] dF_{\mu}(P) \right] dH(\mu) > 0.$$

Hence, the value of an established entrepreneur, $W(S, N)$, is given by the present value of his *current* projects plus the present value of his wage. Future projects have zero value after paying a competitive price for the idea. Similarly, the value of a worker, W_0 , is given by the present value of his wage plus the present value of having ideas. The latter includes the gains from selling some of these ideas and of the possibility of spinning off with one of them.

The key insight in the previous proposition is that the selection of the ideas implemented in existing firms, which is given by μ_L and μ_H , is independent of S and N . Because of this, the set of ideas that will be implemented within each firm is independent of the firm’s size. It is the competitive market for ideas that leads to this result. In the absence of this competitive market, entrepreneurs of existing firms will appropriate some of the surplus of a given idea. Because of this, established entrepreneurs would care about the number of ideas they implement in the future, which in general depends on their size. Hence, selection of ideas and projects would depend on firm size as well. This in turn would determine how many ideas they buy in equilibrium given their size, namely, $\gamma(\lambda, \cdot)$. This is not the case when entrepreneurs pay the competitive market price Z of an idea. In this case, the expected benefits for all entrepreneurs is zero, the selection thresholds are identical for all firms, and the shape of the function $\gamma(\lambda, \cdot)$ (but not its level θ) is indeterminate and so has to be specified exogenously. This is the sense in which the market for ideas is key to generating scale independence in the selection of ideas.¹²

In order for firm growth to be scale independent not only do we need the thresholds μ_L and μ_H to be independent of the size of the firm, but we also need the number of innovations bought by a given firm in the market for ideas to be proportional to its size. In order for firms to buy projects at a rate proportional to their size, we will assume a linear probability of learning about ideas, $\gamma(\lambda, N) = \theta\lambda N$, where θ is determined by equalizing the demand and supply of ideas. In effect, we are assuming that an entrepreneur learns about ideas for sale through all the agents working in his firm, including himself. This may be the result of workers having ideas themselves and selling them to their employer, as that may be easier than contracting with an unknown entrepreneur. Alternatively, one may think of employees finding out about ideas for sale in the market and informing their employer.

Figure 1 summarizes what we have learned about firm behavior and firm entry. It shows expected revenue μ in the real line. For projects above μ_H , firms’ workers spin off. All other

¹² In the particular case in which $\gamma(\lambda, N) = \tilde{\lambda}(\lambda)N$ and utility is linear, even without the market for ideas, μ_L is independent of size. The market for ideas is necessary to obtain scale independence in the selection of ideas if the utility function is not linear and/or $\gamma(\lambda, \cdot)$ is not linear in N . In Section A of the Appendix, where we consider the case of exponential utility, this is evident.

firms implement ideas with μ s between μ_L and μ_H . An incumbent risk-neutral entrepreneur would implement projects as long as they pay expected return w , so the difference between μ_L and w is the result of an entrepreneur's ability to drop the project next period (this ordering can change once we consider risk-averse agents). The threshold that determines μ_L is given by

$$(5) \quad \int (P - w) dF_{\mu_L}(P) + \frac{\beta}{1 - \beta} \int_w (P - w) dF_{\mu_L}(P) = 0.$$

The difference between μ_L and μ_H is the result of the option value of exiting and setting up a new firm in the future, f_0 . The threshold μ_H is implicitly determined by

$$(6) \quad \int (P - w) dF_{\mu_H}(P) + \frac{\beta}{1 - \beta} \int_{P_H} (P - w - f_0) dF_{\mu_H}(P) = Z.$$

The value of having ideas for the worker is given by

$$(7) \quad f_0 = \lambda \int \max \left[\int (P - w) F_{\mu}(P) + \frac{\beta}{1 - \beta} \int \max [P - w - f_0, 0] dF_{\mu}(P), 0 \right] dH(\mu) + (1 - \lambda)Z.$$

This value includes the option to spin off and start a new firm and therefore give away the chance to spin off with a better idea as well as the expected value of the ideas sold in the market λZ . The difference between the two thresholds comes from the option value, included in f_0 , of closing a new firm and starting another one later on with a better idea and the fact that workers give up the price of an idea Z when they set up the firm. Note that if $f_0 = 0$, the two threshold equations and the equation for Z imply that $Z = 0$ and $\mu_L = \mu_H$. So, all projects would be implemented via spinoffs. However, as shown in the previous proposition, f_0 is positive, since workers can extract the value of very good projects by spinning off. New firms will require a higher return from their first project than existing firms demand from new projects, given the larger option value f_0 that new firms have of returning to an old firm and spinning off in the future, namely, $P_H > w = P_L$.¹³ Both of these equations imply that the number of entrants as a fraction of the population is constant and so is the number of new projects implemented in existing firms each period as a fraction of total population.¹⁴

These results depend on the assumption that project returns are specific to the entrepreneur who implements the idea. If project returns were not entrepreneur specific, then a worker who found that his project has a payoff between w and P_H could sell that project to an established firm. Then, workers will not be any more selective in implementing ideas than established firms and $\mu_H = \mu_L$, but it would still be the case that $P_H > w = P_L$.

Finding an equilibrium amounts to solving Equations (4)–(7) for the values of Z , f_0 , μ_L , and μ_H and the market clearing condition (3) for θ . The next proposition establishes that an equilibrium exists and is unique.

PROPOSITION 2. *A long-run equilibrium for this economy exists and is unique.*

¹³ See Section B of the Appendix, where the difference between Equations (5) and (6) is not f_0 but an exogenous entry cost f .

¹⁴ Note that $P_H > w$ implies that a project dropped by a spinoff could be sold to existing firms. Since we have assumed that projects are specific to the entrepreneur, we have abstracted from these transactions. We could relax this by assuming that projects can be sold at some cost. This would lead to a situation where some projects are sold at market value to existing firms. If we assume that projects can be sold at some cost, the agent who is deciding whether to carry on as an entrepreneur will decide to do so if the cost of selling the existing project is greater than the value of being a worker and having the option to spin off in the future. Then, this model would have *both* a market for ideas and a market for start-ups. Importantly, it will not eliminate the market of ideas, which is the focus of this article.

As we noted above, one potential issue is the existence of an equilibrium with contingent contracts, namely, a contract in which an entrepreneur offers the worker the contingent return on an idea minus w . We can now be more specific about how big these costs of writing contingent contracts need to be for such contracts not to be used. Workers with good ideas that would otherwise spin off would be willing to stay if the cost of writing a payoff-contingent contract that mimics what the worker gets by spinning off costs less than f_0 . Note that this includes everything the worker gives up by spinning off, namely, the market price of his future ideas as well as the option to spin off in the future with a better idea.

In Section A of the Appendix, we show that all results, except $\mu_L < w$, hold when individuals order preferences according to $u(c_t) = -ae^{-bc_t}$. In this case all agents in the economy *are* risk averse. However, because their risk aversion does not depend on the level of their wealth, in the presence of markets for ideas, we still obtain the result that μ_L is constant and therefore that the selection of ideas and projects is scale independent as before. All other details, including an explicit formula for the equilibrium price of an idea, are relegated to Section A of the Appendix.

4. COMPARATIVE STATICS, EMPIRICAL IMPLICATIONS, AND EVIDENCE

Our theory determines the selection thresholds μ_L and μ_H , the market price of ideas Z , and the option value of spinning off f_0 . These equilibrium values depend on the different parameters of the model, namely, the discount factor β , the probability of having ideas λ , the outside wage w , as well as the distributions of realizations given an idea, F_μ , and the distribution of the quality of ideas, H . With linear utility, no other parameters enter the model. Given the values of (μ_L, μ_H, Z, f_0) the model also yields implications on the size distribution of firms, firm growth rate, the number of spinoffs, and other distributional outcomes. Important for our purposes is that none of the latter implications have to be solved for in order to obtain (μ_L, μ_H, Z, f_0) . We can solve the model sequentially because the value of ideas, both within existing firms and in new firms, does not depend in equilibrium on the distribution of firm sizes in the economy. We turn to the implications on the size distribution of firms and other outcomes in the next section.

The values of (μ_L, μ_H, Z, f_0) are determined by Equations (4)–(7). From Equation (5) it is immediate that μ_L is decreasing in β , independent of λ , and increasing in w . Furthermore, μ_L decreases if we switch from F_μ to F'_μ where F'_μ first order stochastically dominates F_μ . μ_L is independent of the distribution H . Hence, firms are more selective as wages increase, but less selective if agents are more patient or if the distribution of realization of a given idea improves.

The effect of parameters on the other equilibrium values is much more complex, as the system is only block recursive for μ_L . We therefore proceed with numerical simulations. In all numerical simulations we let $F_\mu(P) = \frac{P - (\mu - \rho/2)}{\rho}$. Namely, we let F_μ be a uniform distribution with range of length ρ centered at μ . For H we use a generalized Pareto distribution with minimum value given by 1 and shape coefficient given by σ . That is,

$$H(\mu) = 1 - (1 + \sigma(\mu - 1))^{-1/\sigma}.$$

So a higher σ implies that the distribution has a left tail with more mass.

Given these two distributional assumptions, we need to pick five parameters. We let $\beta = 0.95$, a standard value for yearly data. We will show comparative statics for the values of ω , λ , and σ . We let $\rho = 8$, which, given the ranges of the other three parameters, gives us realistic ratios of new employees in new firms relative to new employees in continuing firms. This statistic is the only moment that matters to determine the size distribution of firms and other distributional outcomes, as we argue in the next section. We denote this statistic λ_H/λ_L and calculate it to be between 0.07 and 0.12 in the U.S. economy from 1989 to 2003 (we discuss the details of this calculation in the next section).

Figure 2 shows the value of (μ_L, μ_H, Z, f_0) as a function of λ and σ . We choose a range of λ that makes a worker have between 0.2% and 0.6% probability of having an idea each period.

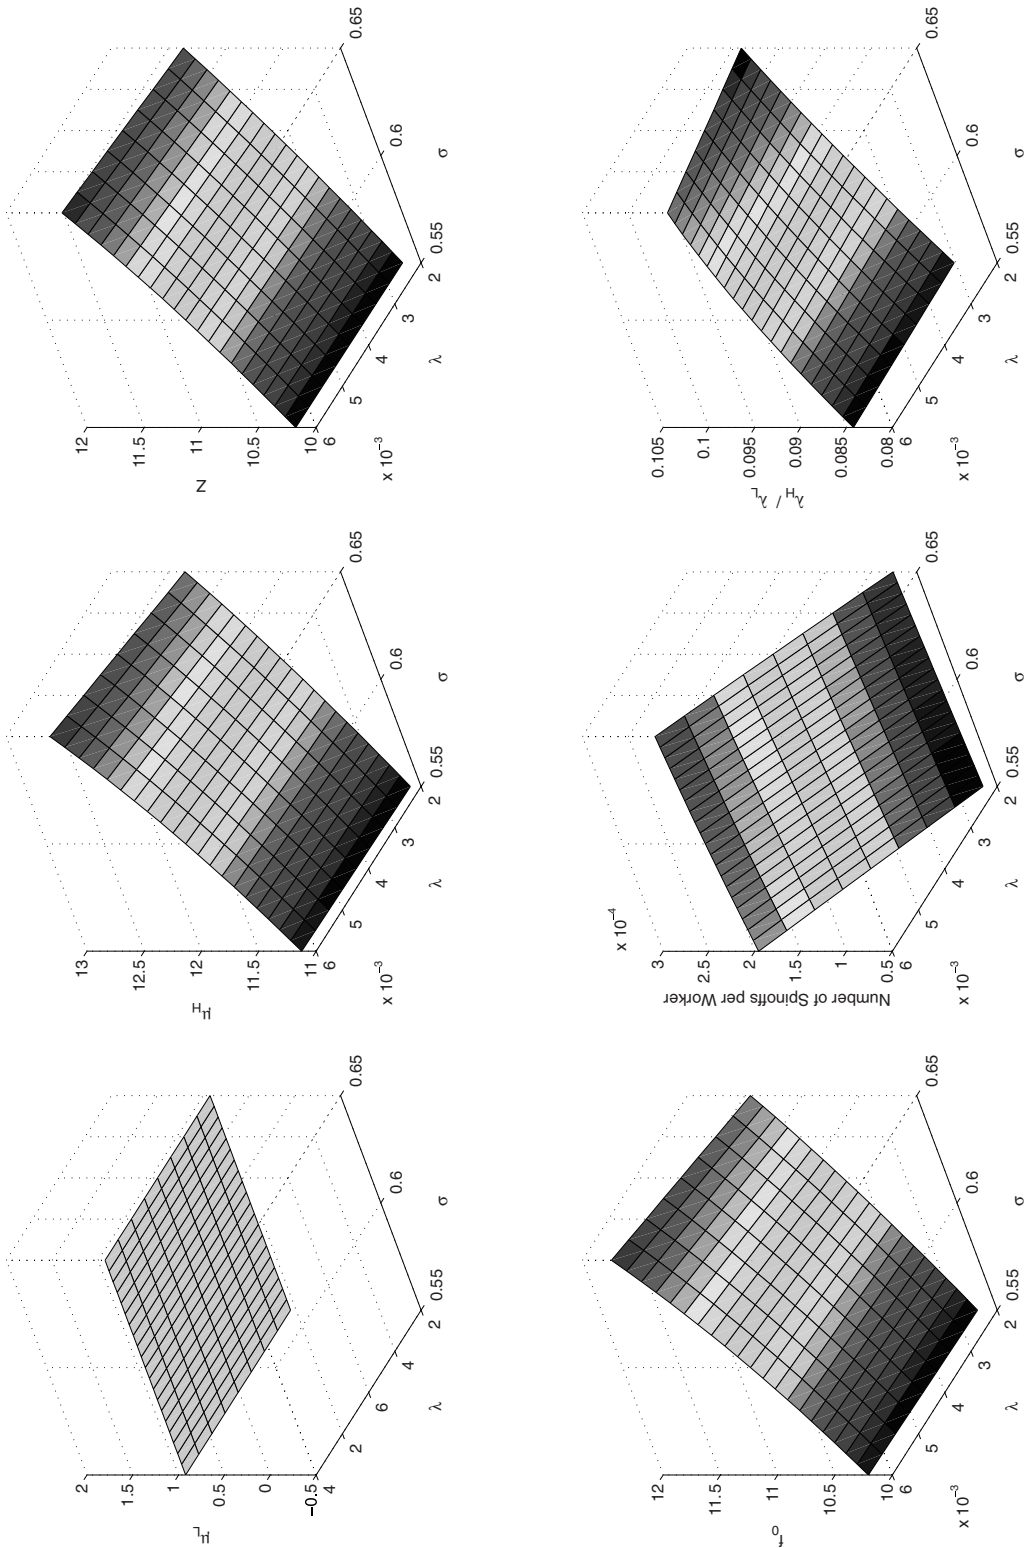


FIGURE 2
COMPARATIVE STATICS FOR λ AND σ




FIGURE 3

THE EFFECT OF λ

So a firm with 1,000 employees will have two to six ideas per year. The range of values of σ that we choose makes agents spin off with about 3% of those ideas. They also show the number of spinoffs per period per worker and the λ_H/λ_L ratio discussed above. As one can see in the figure, the ratio of new employees in new firms to new employees in continuing firms is in the relevant range for these parameter values.

Several results are noteworthy. As σ increases, we shift more mass to the tail of the distribution H . This implies that good ideas are more likely. As noted above, this does not affect how selective continuing firms are in choosing which projects to implement, μ_L , but it does imply that workers wait longer for better ideas to spin off. As σ increases, we also find that the price of ideas increases, since the average idea sold in the market is now of better quality. This is because both μ_H went up and there is more mass in the left tail of the distribution. The option value of setting up a new firm, f_0 , also goes up.

An increase in λ has similar effects, although (in this exercise) smaller in magnitude. First, as λ increases, the values of μ_H , Z , and f_0 all go up. This means that a higher probability of generating ideas leads to more selection of ideas by potential entrepreneurs. Since firms' selection of implemented projects μ_L does not vary with λ , this means that firms implement a wider range of projects ($\mu_H - \mu_L$ grows). Note that even though workers are more selective in their choice of projects (higher μ_H) there are more spinoffs per worker. The effect of λ is hard to assess in Figure 2, so we show this dimension separately in Figure 3.

Figure 4 shows the same six graphs presented in Figure 2, but we change the axes to reflect changes in the wage w instead of the shape parameter of the Generalized Pareto distribution, σ . The figure shows similar comparative statics for λ as described above. The effect of higher outside wages follows a different pattern. Higher w implies more selection by firms, as is easy to show analytically (higher μ_L). Furthermore, a higher w reduces μ_H , Z , and f_0 , but it increases the number of spinoffs and the ratio of lambdas. So if workers are more expensive, firms will implement a smaller range of projects. Since on average the ideas sold to firms are worse, the price of ideas falls, as does the option value of an idea. As the threshold μ_H decreases with the increase in w , more ideas lead to spinoffs, and, given the number of ideas per person λ , there are more spinoffs per person. This also implies that the number of new employees in new firms

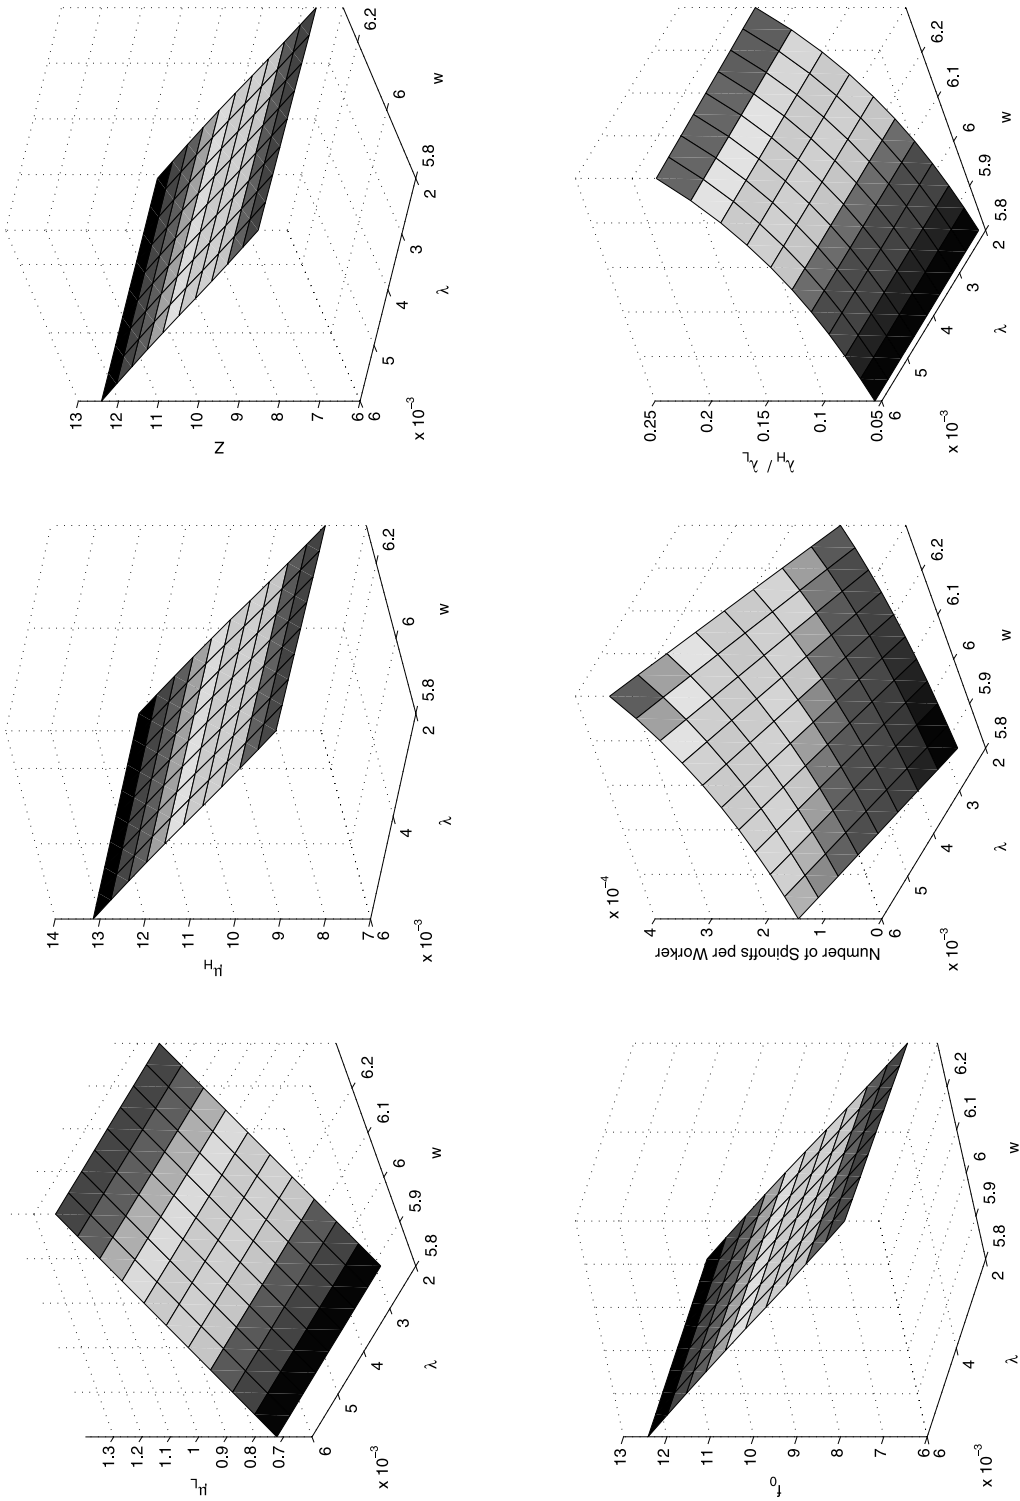


FIGURE 4
COMPARATIVE STATICS FOR λ AND w

grows relative to the number of new employees in continuing firms and so λ_H/λ_L increases as well.

We now discuss the evidence in support of the implications of our theory. The empirical literature on spinoffs has identified several regularities. An important regularity seems to be that “employees start their own firms after becoming frustrated with their employers. Their frustration is often related to having an idea about an innovation or a new (sub)market rejected by their employer” (Klepper and Sleeper, 2005). In our article, disagreements mean that parties cannot find a price at which they can transact. This is the core implication of our way of modeling incomplete information.

Perhaps the most important implication of our framework is that only the best ideas lead to spinoffs. Therefore, we should observe in the cross section that the first idea of a firm is in general better than future ideas. This is consistent with some of the available evidence, which suggests that the first product of a firm is, on average, the most successful of its products. Prusa and Schmitz (1994) argue that this is the case in the PC software industry. The first product of a firm sells, on average, 1.86 times the mean product in its cohort, whereas the second product sells only 0.91 times the mean product in its cohort. That is, first products are, on average, about twice as successful as second products. The first product is also about twice as successful as the third, fourth, and fifth products. This evidence suggests that spinoffs discriminate more than incumbent firms in choosing which projects to implement. This is exactly in line with the selection mechanism our theory underscores. Another related finding is the evidence in Luttmer (2008) that new firms need to draw from a better distribution in order to explain the *age* distribution of firms.

The model also predicts the fraction of unsuccessful spinoffs that exit the economy. Large firms can have unsuccessful projects, too, but they do not exit; they just drop the project. They do not exit, since they have at least one ongoing project that provides a permanent source of profits. Some authors (for instance, Hall and Woodward, 2007) have argued that a common phenomenon is for workers to spin off only to be acquired by a larger firm some years later. Our theory provides a rationale, namely, adverse selection, for why we have spinoffs, but we do not address the issue of spinoffs being acquired by existing firms later on. The theory also has predictions for the entry process of firms. Large firms generate more spinoffs than small firms, although as a fraction of the workforce, the number of spinoffs is constant. This is consistent with the evidence discussed in Klepper and Sleeper (2005) that firms that produce a wider range of products generated more spinoffs over time. Franco and Filson (2006) show, for the hard-drive industry, that more know-how (which is likely correlated with size) also leads to more spinoffs.

Our numerical simulations in the previous subsection generated several other empirical implications. In particular, parameters that could be related to better economic circumstances, such as λ or σ , imply a higher threshold to spin off, more spinoffs, a higher price of ideas, a higher option value of spinning off, and proportionately more new workers entering through spinoffs. These predictions of the model are particularly relevant in light of the empirical facts in Jovanovic and Rousseau (2009). They document that the aggregate Tobin's Q is positively related to the skill premium, negatively related to the relative investment of incumbents, and positively related to the number of spinoffs. High aggregate Tobin's Q can be compared to times in which λ is high: times where many agents have ideas. Or times where σ is high: times where ideas are particularly good. Given this, our model predicts that in fact we should see more spinoffs, a higher value of the option to spin off, a higher market price of ideas—which could be compared to the skill premium where the skilled are the agents with ideas—and less investment by incumbents relative to new firms, as is evident by the increase in λ_H/λ_L .

Similarly, the previous section illustrates the predictions of the model for how these variables change with respect to outside wages. In particular, the model implies that the price of ideas decreases with wages, the number of spinoffs increases, and the entry of new workers through new firms increases as well. We do not know of empirical work that has contrasted these types of implications with data but these are implications of our model that can be potentially tested.

As we show in the remainder of the article, the implications of all of these parameters for the ratio of λ_H/λ_L are particularly relevant, as they are directly related to the shape of the size distribution of firms. We now turn to these implications of the model and examine how they compare with the data.

5. INVARIANT DISTRIBUTION OF FIRM SIZES: THEORY AND EVIDENCE

In order to derive the implications of our model for firm growth and the size distribution of firms, we need to take a stand on the number of projects that firms find out about. The reason is that in our model, given the price in the competitive market for ideas, entrepreneurs are indifferent about how many ideas to buy. This implies, as argued above, that we need to specify the function $\tilde{\gamma}(\cdot)$. Assume that entrepreneurs encounter ideas in proportion to the number of agents in the firm. This amounts to assuming that they find out about ideas generated by agents in their own firm or, alternatively, they and their workers get information about ideas at a constant rate per person.

Suppose that a firm with N projects has a probability of finding out about an idea given by $\tilde{\gamma}(\lambda, N) = \lambda N$. We assume that the maximum size of a firm is given by \bar{N} such that $\lambda \bar{N} < 1$.¹⁵ Everyone in the firm has a probability $\lambda > 0$ of generating an idea. Note that since the value of λ depends on our definition of a period, we can always make λ small enough by appropriately defining the length of a period in the model. Correspondingly, we can make \bar{N} arbitrarily large. In case a firm hits the size constraint \bar{N} , its workers will sell ideas to other firms. For the moment we abstract from this problem, but we return to it below.

One way to justify our choice of $\gamma(\lambda, N) = \theta \lambda N$ is to assume that an entrepreneur can only evaluate one idea per period. Then, if he confronts ideas from each of his workers with probability $\theta \lambda$, the probability of evaluating an idea in any period is multinomially distributed and given by $1 - (1 - \theta \lambda)^N$. Taking a linear approximation at $\lambda = 0$ yields that the probability of an entrepreneur having access to an idea is given by $\theta \lambda N$. Of course, this approximation works well when we are far away from the upper bound $\bar{N} = 1/\lambda$. Otherwise, we would need to work directly with $1 - (1 - \theta \lambda)^N$, which would complicate the analysis but would avoid having to impose an arbitrary upper bound.

As we have shown, entrepreneurs are indifferent about how many ideas to buy in the market. Thus, in combination, this specification of $\tilde{\gamma}(\cdot)$ implies that the growth of firms will be independent of size (Gibrat's Law). It is important to note, however, that Gibrat's Law is not sufficient to determine the form of the invariant distribution of firm sizes. The latter depends on the process of entry and exit. For example, as Gabaix (2005) shows, Gibrat's Law with no entry and exit and a reflecting barrier arbitrarily close to size 0 leads to a Pareto distribution with coefficient 1. In contrast, in Eeckhout (2004) Gibrat's Law with no entry (of cities) and no reflecting barrier leads to a log normal distribution in the limit. In our case, the total mass of firms is also normalized to 1—which is equivalent to having exit at a proportional rate—but there will be entry at size 1. This, as we show below, leads to a Yule distribution for firm sizes, which fits the firm size data well (as shown before by Simon and Bonini, 1958). In this sense, it is the entry process that distinguishes our theory from other theories of firm dynamics that are also consistent with Gibrat's Law.

Given our specification for $\tilde{\gamma}(\cdot)$ we need to determine θ such that the supply of ideas and the demand for ideas equalize in equilibrium at price Z . Let δ_N denote the share of firms of size N in equilibrium (we will discuss this distribution in much more detail below). Then, market

¹⁵ Alternatively we could work with continuous time and assume that the process by which firms generate ideas is Poisson with parameter λN . This would imply an identical random process for generating ideas in continuous time. Note that we are assuming that the process of generating ideas and the process of assigning knowledge and ownership are independent. If instead each worker had an unconditional probability of generating an idea λ independently of other workers, there would be a positive probability of generating several ideas per period, which we rule out.

clearing in the market for ideas requires

$$(8) \quad \lambda \sum_{N=1}^{\infty} (N-1) \delta_N = \theta \lambda \sum_{N=1}^{\infty} N \delta_N,$$

where the l.h.s. is the supply of ideas ($N-1$ workers in a firm of size N have a probability λ of generating an idea) and the r.h.s. is the demand for ideas (a firm of size N learns of [and buys?] $\nu(\lambda, N) = \theta \lambda N$ ideas). Denote by ν the average firm size; then

$$(9) \quad \theta = 1 - \frac{1}{\nu}.$$

That is, the number of ideas per worker encountered by an average entrepreneur is equal to $\lambda \theta$. Each worker has a probability λ of having an idea. Since entrepreneurs do not have ideas, on average each entrepreneur is confronted with a probability of encountering an idea that is lower than λ , namely, $\theta \lambda$. Thus, intuitively, θ is just unity minus the share of entrepreneurs in the total number of people working in the industry ($1/\nu$).¹⁶

Note that our assumption that the maximum size of firms is given by \bar{N} implies an additional adjustment for θ . Since firms get zero expected benefits out of implementing ideas, entrepreneurs are indifferent about expanding or not. Hence, they do not care about this upper bound for the size of their firm. The only role that this bound plays is to determine how other firms grow if there is a positive mass of constrained firms. Thus, the only adjustment we need to make is to add the upper bound \bar{N} to Equation (8). Notice, however, that Equation (9) still holds.

In order to derive the size distribution of firms, first note that the size of the industry will increase constantly in our setup since innovation does not stop (every worker in the industry has probability λ of having an idea independent of where she works). The probability of firms adding a project is positive for all firms, whereas the probability of dropping a project that is already being used is zero. Hence, firms will only grow over time. This is combined with a positive mass of new entrants with one worker every period. So we can show only that there is an invariant distribution of employment shares and firm sizes measured as a share of total employment. That is, we normalize by the size of total employment. This normalization is equivalent to having an exogenous death rate independent of size.¹⁷

First, consider the transition equation for a firm with N workers. Each worker has a probability λ of having an idea. If she does, the firm implements it if $\mu \in [\mu_H, \mu_L(N)]$ and if it implements it, the firm uses the idea with probability $1 - F_\mu(P_L(N))$ where $P_L(N)$ is such that

$$W\left(\frac{NS + P_L(N)}{N+1}, N+1\right) = W(S, N),$$

which by the arguments above does not depend on S .

In what follows we will ignore the upper bound on firm sizes \bar{N} . We will return to it once we define the invariant distribution of firm sizes for the case without this bound. Hence, if

¹⁶ Since entrepreneurs do not get ideas, a new firm with only one worker will need to buy an idea from a nonemployee to grow. More generally, small firms will need to buy ideas from nonemployees in order to grow at the same rate as large firms. The number of ideas originating in a firm that are sold to outside entrepreneurs, as a proportion of the number of workers in the firm, is given by $\lambda[(1-1/N) - \theta] = \lambda[1/\nu - 1/N]$.

¹⁷ In our model, firms die for two different reasons. First, new firms may find out that their productivity is low and exit. Second, as described in the text, we normalize the total mass of firms to 1. This is equivalent to assuming that exit rates of continuing firms are independent of size. Thus, in a stylized way, we do incorporate the feature that exit rates are decreasing in size.

$p(N, N + 1)$ denotes the probability of a firm transitioning from N to $N + 1$ workers

$$p(N, N + 1) = \theta\lambda N \int_{\mu_L}^{\mu_H} (1 - F_\mu(P_L)) dH(\mu).$$

Hence,

$$p(N, N') = \begin{cases} 0 & \text{for } N' > N + 1 \\ \theta\lambda N \int_{\mu_L}^{\mu_H} (1 - F_\mu(P_L)) dH(\mu) & \text{for } N' = N + 1 \\ 1 - \left[\theta\lambda N \int_{\mu_L}^{\mu_H} (1 - F_\mu(P_L)) dH(\mu) \right] & \text{for } N' = N \\ 0 & \text{for } N' < N \end{cases}.$$

Let $S = \{1, 2, \dots\}$; then, for any $A \subseteq S$,

$$p(N, A) = \sum_{N' \in A} p(N, N')$$

is positive if $N \in A$ or $N + 1 \in A$.

Let L_t be the total labor force and E_t the total number of firms or enterprises in period t , and let $\{\delta_N\}$ be the invariant distribution of firm sizes. The probability that a firm with N employees generates a spinoff is given by

$$s(N) = \theta\lambda N \int_{\mu_H} (1 - F_\mu(P_H)) dH(\mu),$$

where P_H satisfies

$$W(P_H, 1) = W_0.$$

Hence, the expected number of spinoffs in period $t + 1$ given the distribution of firm sizes in period t is given by

$$\begin{aligned} E_t \sum_{N=1}^{N=\infty} s(N) \delta_N &= E_t \theta\lambda \int_{\mu_H} (1 - F_\mu(P_H)) dH(\mu) \sum_{N=1}^{N=\infty} N \delta_N \\ &= \theta\lambda \int_{\mu_H} (1 - F_\mu(P_H)) dH(\mu) L_t \equiv \lambda_H L_t, \end{aligned}$$

where λ_H denotes the number of new employees in new firms as a fraction of total employment. Hence the expected number of spinoffs is a constant fraction of the population, L_t .

Similarly, we can calculate the expected number of new workers in existing firms, which is given by

$$\begin{aligned} E_t \sum_{N=1}^{L_t} p(N, N + 1) \delta_N &= E_t \sum_{N=1}^{L_t} \theta\lambda N \int_{\mu_L}^{\mu_H} (1 - F_\mu(P_L)) dH(\mu) \delta_N \\ &\equiv \lambda_L E_t \sum_{N=1}^{L_t} N \delta_N = \lambda_L L_t, \end{aligned}$$

where λ_L denotes the number of new employees in old firms as a fraction of total employment. Then, for E_t large

$$L_{t+1} = L_t + E_t \sum_{N=1}^{L_t} [p(N, N+1) + s(N)] \delta_N = (1 + \lambda_H) L_t + \lambda_L L_t.$$

Given our definition of λ_L and λ_H , population evolves according to

$$L_{t+1} = (1 + \lambda_H + \lambda_L) L_t.$$

Thus,

$$E_{t+1} = E_t + E_t \sum_{N=1}^{\infty} s(N) \delta_N = E_t + \lambda_H L_t.$$

Hence the number of firms is expanding at a constant rate. In terms of the number of firms, the economy is growing at a constant rate. Note that we are assuming that E_t is large enough so that L_t and E_t evolve deterministically. For small L_t and E_t , however, both are random variables that evolve according to a stochastic process.

We now compute the invariant distribution of the share of workers in firms of different sizes. Let ϕ_N denote the probability that a worker is employed by a firm with N workers. The probability that a worker has an idea that is used within the firm is given by λ_L , independently of the firm's size. Then, the invariant distribution satisfies

$$[\phi_1 (1 - \lambda_L) + \lambda_H] L = \phi_1 L' = \phi_1 (1 + \lambda_L + \lambda_H) L$$

or

$$\phi_1 (1 - \lambda_L) + \lambda_H = \phi_1 (1 + \lambda_L + \lambda_H),$$

which implies

$$(10) \quad \phi_1 = \frac{\lambda_H}{\lambda_H + 2\lambda_L}.$$

Intuitively, the number of workers in firms of size 1 today, $\phi_1 L$, minus the number of workers in firms of size 1 that become workers in firms of size 2, $\phi_1 \lambda_L L$, plus the number of new workers in firms of size 1, $\lambda_H L$, is equal (in the invariant distribution) to the number of workers in firms of size 1 tomorrow, $\phi_1 L'$, which is equal to $\phi_1 (1 + \lambda_L + \lambda_H) L$, given that the growth rate of employment is $\lambda_L + \lambda_H$.

Similarly, for firms of size N ,

$$\begin{aligned} \phi_N (1 - \lambda_L N) + \phi_{N-1} \lambda_L (N-1) + \phi_{N-1} \lambda_L &= \phi_N (1 - \phi_L N) + \phi_{N-1} \lambda_L N \\ &= \phi_N (1 + \lambda_L + \lambda_H) \end{aligned}$$

and so

$$(11) \quad \phi_N = \phi_{N-1} \frac{\lambda_L N}{\lambda_H + \lambda_L (N+1)},$$

which implies that

$$\frac{\phi_N}{N} = \frac{\phi_{N-1}}{N-1} \frac{\lambda_L (N-1)}{\lambda_H + \lambda_L (N+1)}.$$

Note that by definition

$$(1 + \lambda_L + \lambda_H) \sum_{N=1}^{\infty} \phi_N = \sum_{N=2}^{\infty} [\phi_N (1 - \lambda_L N) + \phi_{N-1} \lambda_L (N-1) + \phi_{N-1} \lambda_L] + \phi_1 (1 - \lambda_L) + \lambda_H,$$

which implies that

$$(\lambda_L + \lambda_H) \sum_{N=1}^{\infty} \phi_N = \left[\sum_{N=2}^{\infty} \phi_{N-1} \lambda_L + \lambda_H \right] = \lambda_L \sum_{N=1}^{\infty} \phi_N + \lambda_H.$$

Hence,

$$\sum_{N=1}^{\infty} \phi_N = 1,$$

and so the resulting ϕ s form a probability distribution. This distribution is the invariant distribution of employment shares across firms of different sizes.

PROPOSITION 3. *There exists a unique invariant distribution ϕ of employment shares across firm sizes, where ϕ_N denotes the share of workers employed by firms of size N .*

To obtain the distribution of firm sizes we need to transform the distribution of worker shares into a distribution of firm sizes. For this, note that if the share of the population employed by firms of size N is given by ϕ_N , then the share of firms of size N , δ_N , is given by

$$(12) \quad \delta_N = \frac{\phi_N}{N \sum_{N=1}^{\infty} \frac{\phi_N}{N}}.$$

Clearly, since $\sum_{N=1}^{\infty} \phi_N = 1$, $0 < \sum_{N=1}^{\infty} \frac{\phi_N}{N} < 1$, and so δ_N is well defined, exists, and is unique. Since we are normalizing the total mass of the size distribution to 1 we are, in effect, introducing exit at a constant rate for all sizes.

COROLLARY 1. *There exists a unique invariant distribution δ of firm sizes.*

Simon and Bonini (1958) propose an exogenous growth and entry process of firms that leads to the same type of distribution, namely, a Yule distribution. This distribution approximates a Pareto in the upper tail. Note from the previous equations that the distributions ϕ and δ depend only on the value of the ratio λ_H/λ_L . So, one of the contributions of our theory is to provide a microfoundation for the value of λ_H/λ_L that in Simon and Bonini (1958) is just an exogenous parameter.

Note also that in the theory $N \leq \bar{N}$. Hence, in order to get distributions of employment shares and firm sizes that are consistent with the theory we need to renormalize both distributions.

Hence, the distribution of employment shares is given by

$$\tilde{\phi}_N = \frac{\phi_N}{\bar{N}} \frac{1}{\sum_{N=1} \phi_N}$$

and the distribution of firm sizes by

$$\tilde{\delta}_N = \frac{\tilde{\phi}_N}{N \sum_{N=1} \frac{\tilde{\phi}_N}{N}}$$

Now consider the expected growth rate of employment, g_N , of a firm of employment size N . The firm grows by one employee with probability $N\lambda_L$; thus

$$g_N = \frac{(N+1)N\lambda_L + N(1-N\lambda_L) - N}{N} = \lambda_L.$$

Hence, the expected growth rate of firms is just given by the probability per worker of its employees generating an idea that is used. This probability is constant, so the expected growth rate in terms of employees of existing firms is constant, which is a statement of Gibrat's Law. Therefore, the model is consistent with the evidence in Sutton (1997), who argues that the unconditional (on survival) growth rate is consistent with Gibrat's Law. Of course, this is the result of our assumption that $\gamma(\lambda, N) \propto \lambda N$; however, we were free to assume this, given that the market for ideas implies that firms with different characteristics select ideas identically. For our purposes, whether Gibrat's Law holds exactly or not is not crucial. What we will show is that using Gibrat's Law we can approximate the size distribution well. However, we can incorporate any pattern of scale dependence in growth rates and use the model to generate implications for the size distribution.

The model has another relevant implication. Since employment grows by one unit at a time, the variance of the growth rate of employment is decreasing in N . In particular, the variance of the growth rate of employment is given by

$$\left[\frac{1}{N} + 2(1 - \lambda_L) \right] \lambda_L + (1 - \lambda_L)^2,$$

which is decreasing in N . Gabaix (2011) documents that the volatility of firm growth rates decreases with size with an elasticity between 0.15 and 0.20. Our model implies a volatility that also declines with size but not with a constant elasticity.

PROPOSITION 4. *The expected growth rate in employment size of existing firms is independent of their size. Furthermore, the variance of firm employment growth is decreasing in firm size.*

Similarly, the expected growth rate in average revenue of a firm with average revenue S and N employees is given by

$$\begin{aligned}
 g_{S,N} &= \frac{\theta\lambda N \int_{\mu_L}^{\mu_H} \int_{P_L}^{\infty} \left(\frac{NS + P}{N + 1} \right) dF_{\mu}(P) dH(\mu) + S(1 - N\lambda_L) - S}{NS} \\
 &= \frac{\frac{\theta\lambda N}{N + 1} \int_{\mu_L}^{\mu_H} \int_{P_L}^{\infty} PdF_{\mu}(P) dH(\mu) + \left(\frac{NS}{N + 1} \right) N\lambda_L + S(1 - N\lambda_L) - S}{NS} \\
 &= \frac{1}{N + 1} \left[\frac{\theta\lambda}{S} \int_{\mu_L}^{\mu_H} \int_{P_L}^{\infty} PdF_{\mu}(P) dH(\mu) - \lambda_L \right].
 \end{aligned}$$

This implies, as μ_L and P_L , and therefore λ_L , are independent of N , that

$$E_S(g_{S,N}) = 0.$$

So average growth rates across firms of different average revenues are zero. However, large firms that have had good realizations and therefore have a high S will tend to grow more slowly, and vice versa. In this sense there will be reversion to the mean, conditional on the number of employees. This is consistent with Luttmer (2008), who argues that a higher rate of growth for young (and small) firms is key in explaining the *age* distribution of firms. Our article complements Luttmer (2008) by providing a microfounded model of why this is the case.

Also note that the variance of $g_S(S, N)$ is decreasing in N , since the larger the firm, the smaller the contribution of new projects (see also Proposition 4 and Gabaix, 2011). Since firms implement projects that yield only nonnegative profits, this implies that the growth rate of total revenue or total profits will decline with size.

Note that since

$$\phi_N = \phi_{N-1} \frac{\lambda_L N}{\lambda_H + \lambda_L (N + 1)},$$

it is immediate that as $\bar{N} \rightarrow \infty$ or $\lambda \rightarrow 0$, when $N \rightarrow \bar{N}$, $\phi_N \approx \phi_{N-1}$, so the share of workers at large firms is approximately constant. This implies that the density of firm sizes will be proportional to $1/N$ as N becomes large. That is, the tail of the distribution will be arbitrarily close to the tails of a Pareto distribution with coefficient 1. Similarly if λ_H is small, $\phi_N \approx \phi_{N-1}(N/(N + 1))$, and so for N large $\phi_N \approx \phi_{N-1}$ and the distribution of firm sizes is approximately Pareto with coefficient 1. This is interesting given that several authors have concluded that the upper tail of the distribution of firm sizes is close to a Pareto distribution with coefficient 1 (see, for example, Axtell, 2001). We summarize these results in the following proposition:

PROPOSITION 5. *As $\lambda \rightarrow 0$, or $\bar{N} \rightarrow \infty$, the density of firm sizes is arbitrarily close to the density of a Pareto distribution with coefficient 1, for large enough firm sizes. Furthermore, the distribution of firm sizes is closer to a Pareto distribution with coefficient 1, the smaller the mass of workers in new firms, λ_H .*

The invariant distribution of firm sizes, as well as any other outcome of the model, is a function of the exogenous parameters and distributions in the model, namely, β , w , and the distributions F_{μ} and H . However, as we show above, the effect of all those variables can be summarized through the value of λ_L and λ_H/λ_L . Therefore, we can assign a particular value to this ratio and compute the resulting distribution of employment shares and firm sizes. Figure 5 illustrates the invariant distribution in this model and compares it with the distribution of firm sizes in 2000 in the United States. The data for the distribution of enterprises come from the Statistics of

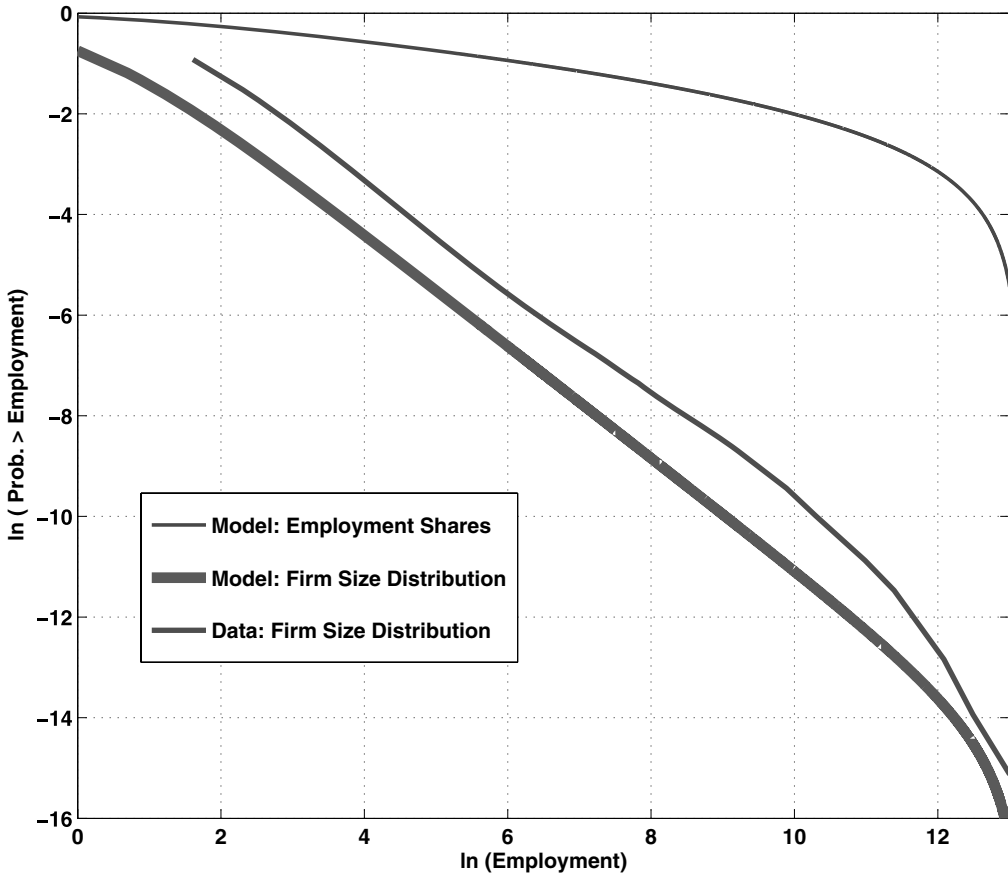


FIGURE 5

MODEL VERSUS DATA

U.S. Businesses (SUSB) program. These data cover the whole private U.S. economy (except agriculture) and were constructed by the U.S. Census Bureau for the study in Rossi-Hansberg and Wright (2007).¹⁸ In this article, we make use of the evidence on enterprises instead of the evidence on establishments, which is the focus of Rossi-Hansberg and Wright (2007). Apart from the size distribution of firms for the aggregate economy, the data set includes the size distribution of enterprises at the two-digit NAICS level for 2000, which we also use below. The industry data are provided for enterprises of up to 10,000 employees.

In order to compute the distribution given in Equation (12), we need to truncate the distribution of firm sizes at a certain size. We choose $\bar{N} = 500,000$, since the largest firms reported in the aggregate data have this number of employees. We choose $\lambda_H/\lambda_L = 1/9$ and so 90% of the new employees are hired by existing firms and 10% by new firms. As is evident from Equations (10) and (11), the distribution depends only on the ratio λ_H/λ_L and not on λ_H and λ_L separately.

Figure 5 shows how our model can do a good job in matching the distribution of firm sizes. Relative to a Pareto distribution with coefficient 1 (a straight line with slope -1), it exhibits a relatively thinner tail of small firms. Furthermore, since in the model and in the data we are truncating the distribution, both distributions exhibit thinner tails than the Pareto distribution for very large sizes. This is only the result of truncation in the model. If we let $\bar{N} \rightarrow \infty$, then

¹⁸ Rossi-Hansberg and Wright (2007) discuss the important differences observed in the size distribution of establishments and enterprises (firms). Their paper is solely concerned with establishments, whereas the current article is about firms.

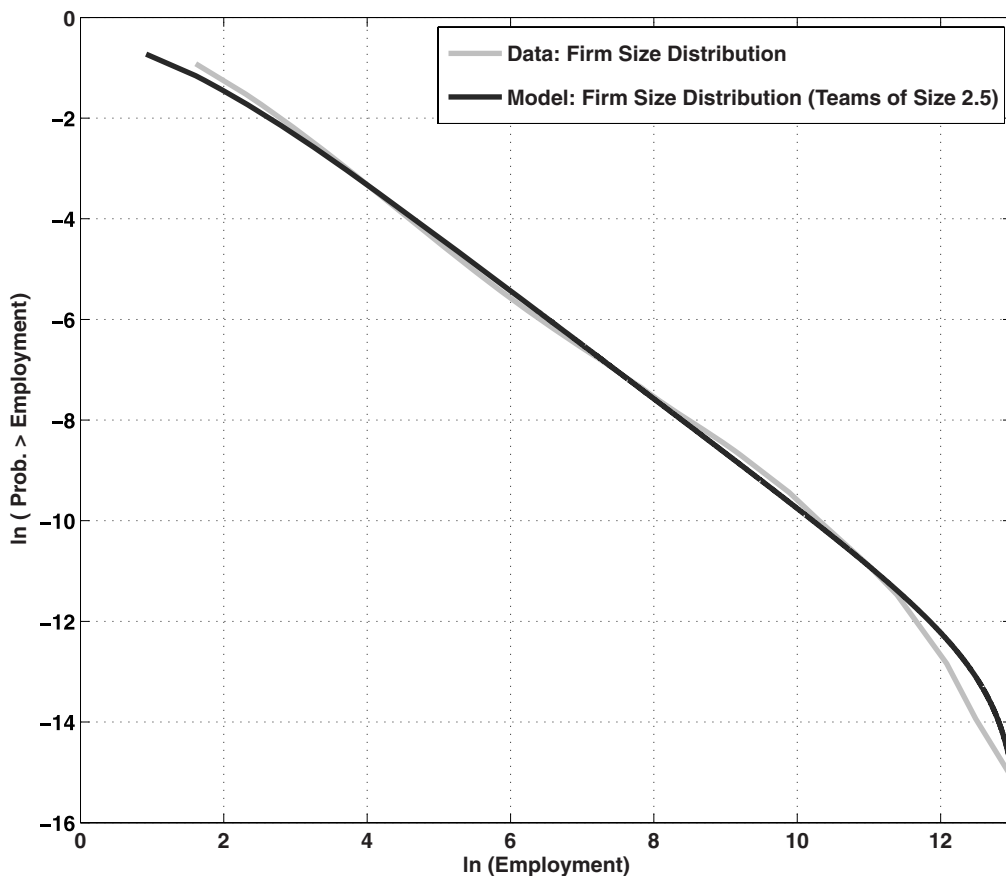


FIGURE 6

MODEL VERSUS DATA WITH VARIABLE TEAM SIZE

the theory implies that the upper tail will be arbitrarily close to the Pareto distribution for large enough sizes. Similarly, in the data the census does not reveal the sizes of the largest firms because of confidentiality concerns. Furthermore, whereas in the model we do not have integer constraints and so there are some firms at all sizes, in the data there cannot be any fractional firms, which truncates the distribution as well. Hence, the reason to have a small mass of large firms on the upper tail is similar in the model and the data.

One noticeable difference between the distribution generated by the model and the distribution generated in the data is that the theoretical distribution lies below the empirical one in Figure 5. The reason for this is that, in our setup, ideas are generated by one employee and not by teams of employees. Consequently, individual agents spin off, not teams of agents. This is clearly not true in the data. Firms enter small, but not necessarily with one employee. Were we to assume that teams of between two and three employees have ideas and spin off, we could shift the theoretical curve in Figure 5 so that it lies on top of the empirical one. Thus, the emphasis is on the shape, not on the level of the curves. This exercise is illustrated in Figure 6 .

Figure 6 uses a value of $\lambda_H/\lambda_L = 0.0736$. This value is the empirical counterpart of λ_H/λ_L in the U.S. economy from 1989 to 2003. Namely, we calculated the net number of workers added by new firms and divided it by the net number of workers added by continuing firms. The data come from the U.S. Small Business Administration (SBA). If we instead calculate the ratio for each year and average over time, λ_H/λ_L increases to 0.1235. Clearly, the model does a very good job of matching the aggregate distribution, arguably better than a Pareto distribution at the upper and lower tails.




FIGURE 7
THE EFFECT OF λ_H/λ_L

Figure 7 shows how we can modify the shape of the invariant distribution by changing the ratio λ_H/λ_L . We illustrate this using $\bar{N} = 10,000$. The relationship between λ_H/λ_L and the primitives of the model is discussed in Section 4. Here, we simply display the invariant distribution for $\lambda_H/\lambda_L = 1/2, 1/5, 1/10, \text{ and } 1/20$. It is clear from the figure that as we increase the number of entrants (by increasing λ_H/λ_L), we shift mass to the lower tail, and, therefore, the slope of the curve in Figure 7 becomes steeper. However, as we know from the previous proposition, if one increases \bar{N} to large enough values, the shape of the distribution approaches a Pareto distribution, as N increases, in all these cases. Figure 7 illustrates the invariant distribution for a large range of parameters. The empirical value of λ_H/λ_L is as calculated above, equal to 0.0736.

The next step is to assess whether the relationship between λ_H/λ_L and the size distribution of firms can explain some of the variation in the size distribution of enterprises across sectors. Different sectors have different values of λ_H/λ_L . This, according to our model, should explain some of the variation in the distribution of enterprise sizes in the data.

There are two main issues in performing this exercise. The first is that the data to calculate λ_H/λ_L for enterprises at the two-digit NAICS level are not publicly available. There are data provided by the U.S. SBA to do this for establishments, but aggregating to enterprises is not trivial. Furthermore, the data for establishments will, in general, show relatively more new net employment in new establishments than in continuing establishments compared to the enterprise data. This occurs since continuing enterprises may add new employees by adding a new establishment. Hence, λ_H/λ_L calculated with establishment data will be larger than with enterprise data. Given the available data, we can compute this measure for the aggregate economy for both enterprises and establishments. In fact, our calculations show that the ratio

for establishments is about 10 times larger than for enterprises. This provides a measure of the bias when we use establishment data instead of enterprise data. We can then compute λ_H/λ_L using establishment data for each two-digit NAICS industry and multiply them by the ratio of the measure of λ_H/λ_L for enterprises divided by the one for establishments for the aggregate economy. This provides us with an estimated measure of λ_H/λ_L for enterprises in all two-digit NAICS industries.

The second issue is that we need a benchmark to compare our results. In particular, we need to assess if the model can explain some of the variation in the observed size distribution of enterprises across industries. We use the aggregate observed distribution of enterprises as our benchmark. So we compare the average absolute distance between the industry size distribution of establishments and the aggregate one in the data. We then compute the size distribution predicted by the model if we use the industry λ_H/λ_L calculated as described above. We need to do one more adjustment. From our discussion of Figures 5 and 6, it is clear that in order to match the level of the size distribution of enterprises we need to adjust the size of the teams that can have an idea and implement it. We saw in Figure 6 that a team of 2.5 workers does a good job for the aggregate economy. However, the required minimum team size probably varies substantially across industries. To address this we choose the team size for each industry so as to minimize the average absolute deviation between the model's implied distribution and the industry distribution of enterprises.

Figures 8 and 9 show the results of this exercise for 19 two-digit NAICS industries and the aggregate economy. For all industries we set the upper bound on firm size, N , equal to 20,000, although our calculations do not change significantly if we choose numbers between 10,000 and 30,000. The data are *censored* at 10,000, so 20,000 seems reasonable.

The figure for the total economy is almost identical to Figure 6, except that the estimated team size is actually 2.72 not 2.5 as in that figure. The figures include the estimated team size for each industry, which varies from 1 in the "Accommodation and Food Services Industry" to 403 for "Management of Companies and Enterprises" (a clear outlier given the shape of the distribution of enterprises in this industry. The second largest team size is 16.4 for "Utilities"). We also report the average absolute difference between the modeled distribution and the data and the aggregate distribution and the industry data. The model beats the aggregate benchmark in 14 out of 19 industries. It does a particularly good job in matching the distribution in "Manufacturing," "Educational Services," and "Health Care and Social Assistance." We conclude that although certainly not the only source of variation, λ_H/λ_L explains part of the industry variation in the size distribution of firms.

As discussed above, in computing the distributions predicted by the model for the observed λ_H/λ_L in each industry, we also estimated the size of the implied team sizes across industries, namely, the size of the team that makes the model fit the industry data as closely as possible. It is then a natural question whether the team sizes implied by the model via the exercises performed above are related to observed "team sizes." Since spinoffs always start with one team, we can look at the average size of entering firms in the data. These data are publicly available from the SBA for firm births between 2002 and 2003, 2003 and 2004, and 2004 and 2005. The data provided by the SBA report the number of employees in new firms by size bins (1–4, 5–9, 10–19, 20–100, 100–500, and 500+). We compute average entry size in two ways. The first one uses the mean employment size in the bin and 750 employees for the 500+ category. The second, a lower bound, is computed by using the lower limit of the bin, which we know exactly (even for the largest bin). We calculate correlations between the model's predicted team sizes as reported in Figures 8 and 9 and the average entry size over the 3 year pairs. The result is a correlation of 0.918 when we use the mean of the size bin and 0.925 when we use the lower limit of the size bin. These very high correlations are, we believe, an encouraging outcome for the model. In particular, note that no data on entry size were used in calculating the model's team size predictions. The results are similar if we use individual year pairs. All correlations fall between 0.909 and 0.936.

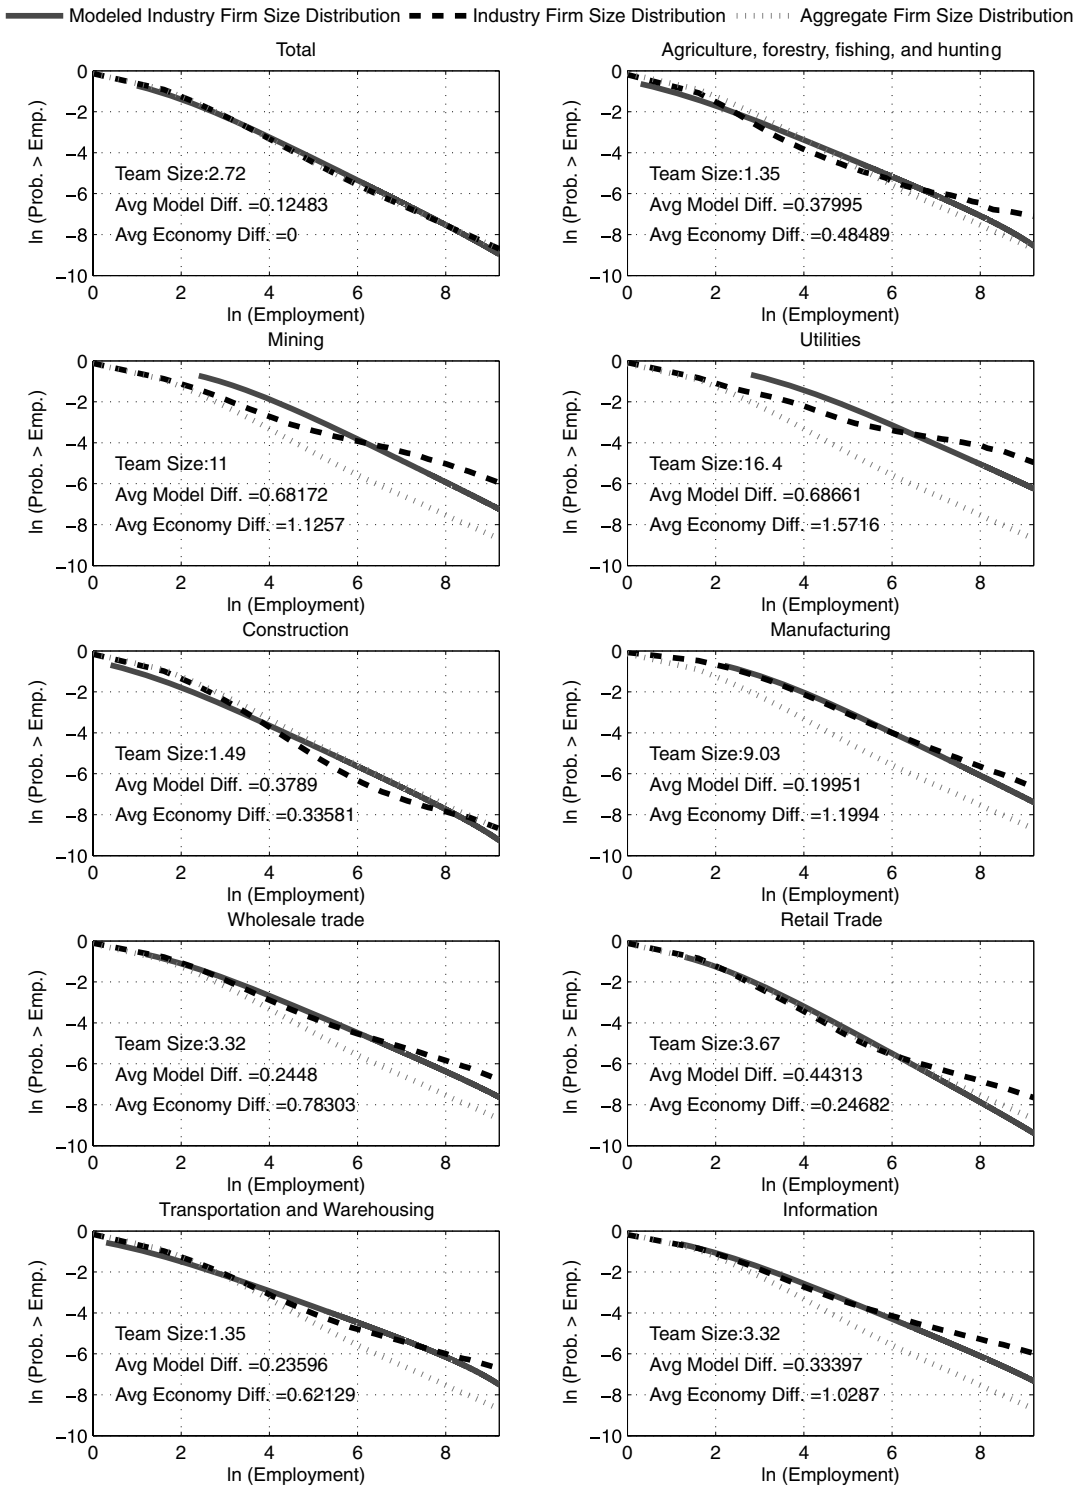


FIGURE 8

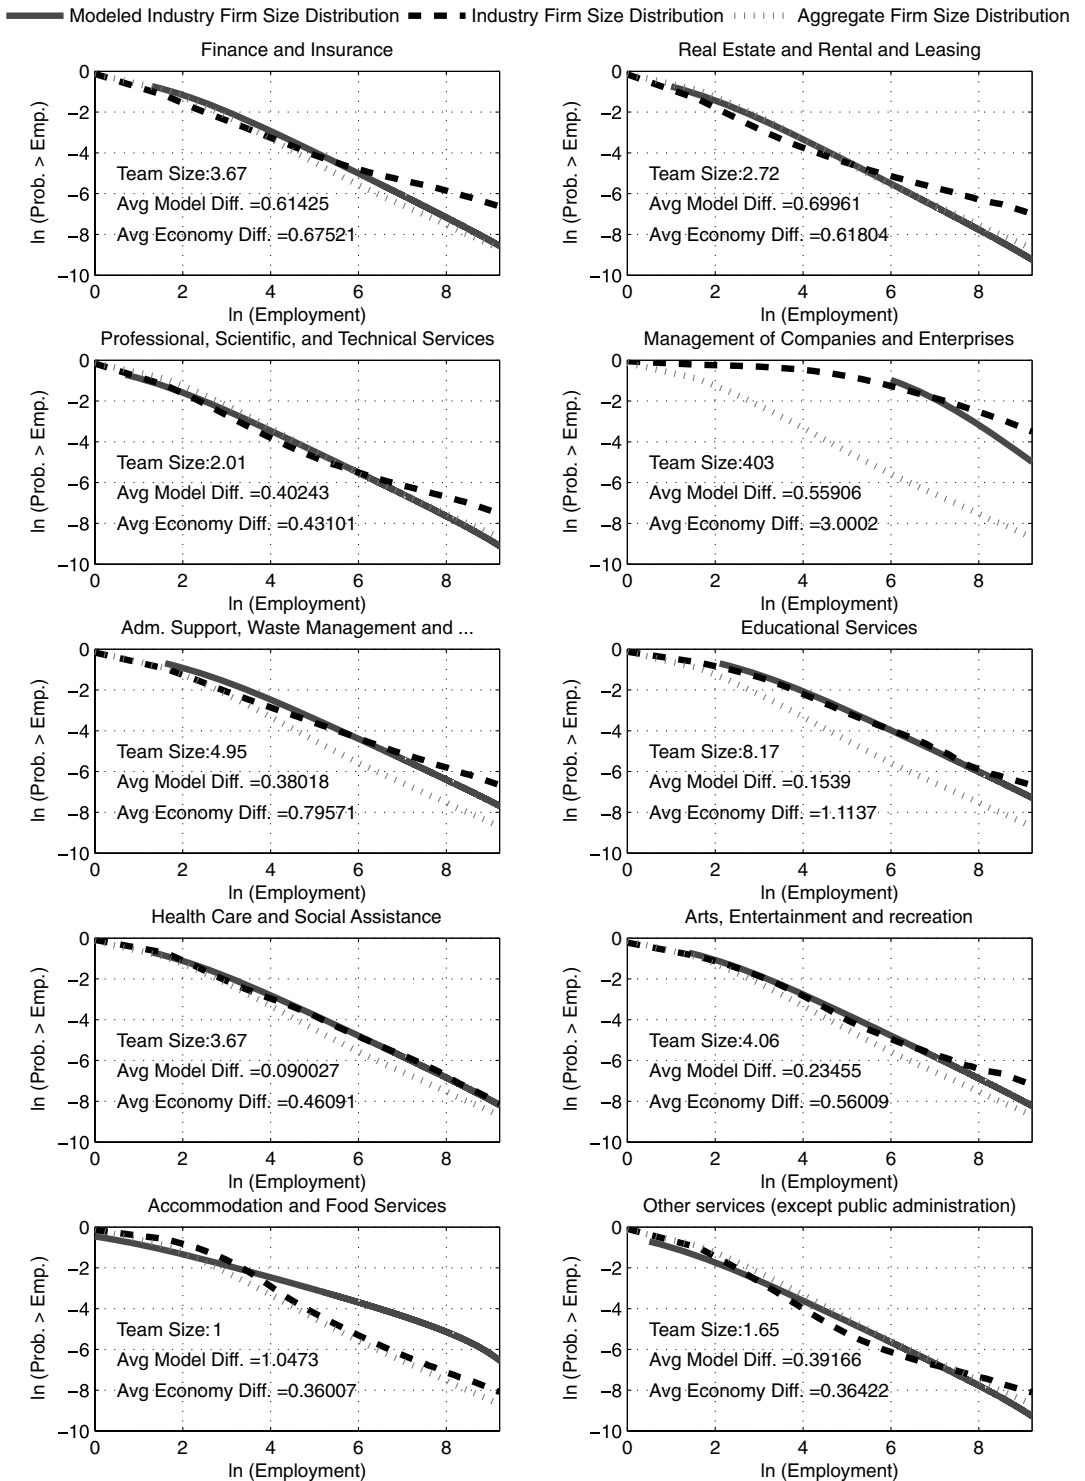


FIGURE 9

6. CONCLUSION

We propose a theory of firm dynamics in which workers have ideas for new projects that can be implemented inside existing firms or, at a cost, in new firms: spinoffs. Workers have private information about the quality of their ideas. Because of an adverse selection problem, workers can sell their ideas to existing firms only at a price that is not contingent on their information. Therefore, workers with very good ideas decide to spin off and set up a new firm. This implies that, through selection, our theory determines the productivity distributions of existing and new firms. Furthermore, since entrepreneurs of existing firms pay a price for the ideas sold in the market that implies zero expected profits, the model is consistent with scale-independent growth in firm size. It is the existence of a competitive market for ideas that is key for this result and, together with entry via spinoffs, leads to distributions of firm sizes that resemble the empirical ones.

The theory produces a size distribution of firms that depends on all the parameters and distributions of the model through the value of λ_H/λ_L only. The model links this ratio to primitive parameters and therefore guides us on how these parameters determine the size distribution of firms. We calculated this ratio using aggregate and industry U.S. data and showed that one obtains size distributions that, for the entire U.S. economy and most industries, are hard to distinguish from their empirical counterpart.

We view as one of the contributions of this article the provision of a simple model of firm dynamics that focuses on private information in the generation of new technology. In order to provide this, we have made some strong assumptions, many of which can be relaxed in different ways. For example, we have assumed that entrepreneurs do not generate ideas, as they are involved in running their firm. In Section B of the Appendix we outline how to relax this assumption by introducing a resource cost of setting up a new firm. We leave other extensions for future research. One extension would be to allow for noncompete clauses. For instance, we could assume that, with some probability, which depends on local laws, workers that spin off may lose property rights to their first project. Gilson (1998) points out that noncompete clauses are not enforced in California but are enforced in Massachusetts. If there is some chance that the parent firm would get the idea, μ_H will increase and make workers more selective about spinning off.

Another extension is to study the links between the aggregate generations of ideas, wages, and industry productivity. We could endogenize wages by requiring that the total compensation package of workers (which includes wages, the price of ideas, and the possibility of getting an idea good enough to become an entrepreneur) clears a fully specified labor market. Then, an increase in the rate at which workers get ideas will have consequences for the observed wages and the measured productivity of firms. Essentially, we could embed this theory of firm dynamics and entry into an equilibrium framework, such as the neoclassical growth model, and study the interactions between industry evolution and wages.

Finally, if spinoffs tend to be geographically close to their parent firm, then they are a potentially important reason why we get clusters of firms working in the same line of business in the same locality. Augmented by a location choice, our theory could then form the basis of a dynamic theory of industrial agglomeration.

APPENDIX

In this appendix we first provide all the proofs not included in the main text. We also prove the analog of Proposition 1 for the case of exponential utility.

A. *Proofs.*

A.1. *Proof of Lemma 1.* Let \mathcal{C} be the space of bounded continuous functions defined on $\mathbb{R} \times \mathbb{N}$. Define the operator $T(W) : \mathcal{C} \rightarrow \mathcal{C}$ as

$$\begin{aligned} T(W)(S, N) &= \gamma(\lambda, N) \int^{\mu_H} \max[V(\mu, S, N) - u(\pi(S, N) - Z + w) - \beta W(S, N), 0] dH(\mu) \\ &\quad + (1 - \gamma(\lambda, N) H(\mu_H)) [u(\pi(S, N) + w) + \beta W(S, N)] \\ &\quad + \gamma(\lambda, N) H(\mu_H) [(\pi(S, N) - Z + w) + \beta W(S, N)], \end{aligned}$$

where

$$\begin{aligned} V(\mu, S, N) &= \int [u(\pi(S, N) + w + P - Z - w)] dF_\mu(P) \\ &\quad + \beta \int \max \left[W \left(\frac{NS + P}{N + 1}, N + 1 \right), W(S, N) \right] dF_\mu(P). \end{aligned}$$

It is easy to show that T is a contraction using Blackwell's conditions. It satisfies monotonicity since $W \leq W'$ implies that $T(W) \leq T(W')$ (all expressions above are increasing in the function W). It satisfies discounting, since for $a > 0$,

$$T(W + a) \leq \beta a + T(W),$$

where $\beta < 1$. Hence, T is a contraction by the Contraction Mapping Theorem, and a unique fixed point to the operator T exists.

Suppose W is increasing in S . Since S appears only in π and W in the definition of T , $T(W)$ is strictly increasing in S , given that π is strictly increasing in S . Hence, since the subset of weakly increasing functions in \mathcal{C} is closed, by the Contraction Mapping Theorem, the fixed point of T is weakly increasing in S . The corollary to the Contraction Mapping Theorem (see page 52, Corollary 1, in Stokey et al., 1989) then says that we can apply the operator T to the fixed point one more time to conclude that the fixed point is strictly increasing in S .

To show that W is continuous in S , note that since the space of continuous functions is closed in the sup norm, we can apply the same argument to show continuity given that π is continuous in S , F_μ is continuous in μ , and $H(\mu)$ is continuous in μ . ■

A.2. *Proof of Lemma 2.* By Lemma 1 $W(S, N)$ exists and therefore $V(\mu, S, N)$ exists. Since $\pi(S, N) + w - w + P - Z$ is strictly increasing in P , the first term in V is strictly increasing in μ , since $\int f(P) dF_\mu(P)$ is increasing in μ for all increasing functions f . Since $W(S, N)$ is strictly increasing in S by the previous lemma,

$$\int \max \left[W \left(\frac{NS + P}{N + 1}, N + 1 \right), W(S, N) \right] dF_\mu(P)$$

is strictly increasing in μ , which proves the result.

That V is strictly increasing in S follows from W being strictly increasing in S .

To show that V is continuous in μ and S , note that by Lemma 1, $W(S, N)$ is continuous in S , the maximum of continuous functions is continuous, and F_μ is continuous with respect to μ . ■

A.3. *Proof of Lemma 3.* First note that given our assumption that $F_\mu(0) = 0$ all μ we know that $\mu \geq 0$. Then

$$V(0, S, N) = u(\pi(S, N) - Z) + \beta W(S, N) < u(\pi(S, N) - Z + w) + \beta W(S, N)$$

since $w > 0$. In contrast,

$$\begin{aligned} \lim_{\mu \rightarrow \infty} V(\mu, S, N) &= \lim_{\mu \rightarrow \infty} \int u(\pi(S, N) + w - Z - w + P) dF_{\mu}(P) \\ &\quad + \beta \int \max \left[W \left(\frac{NS + P}{N + 1}, N + 1 \right), W(S, N) \right] dF_{\mu}(P) \\ &> u(\pi(S, N) - Z + w) + \beta W(S, N). \end{aligned}$$

The inequality follows because the probability that $P > w$ goes to 1 as $\mu \rightarrow \infty$ and

$$\max \left[W \left(\frac{NS + P}{N + 1}, N + 1 \right), W(S, N) \right]$$

is nonnegative. Hence, since $V(\mu, S, N)$ is strictly increasing and continuous in μ , by the Mean Value Theorem there exists a unique scalar μ_L that satisfies Equation (1) for each pair (N, S) . Let $\mu_L(S, N)$ be the unique function that takes this value given a pair (S, N) . ■

A.4. Proof of Lemma 4. We first need to show that V_0 is increasing and continuous in μ , but this follows directly from F_{μ} being continuous and increasing in μ as assumed above. We also know that

$$V_0(0) = u(0) + \beta W_0 < u(w + Z) + \beta W_0$$

since $w + Z > 0$ and $\max[W(0, 1), W_0] = W_0$. The latter is the result of the fact that a new entrepreneur with a project that pays 0 earns less than a worker and has fewer opportunities with regard to exploiting his future ideas (he has exercised the option of spinning off). In contrast,

$$\begin{aligned} \lim_{\mu \rightarrow \infty} V_0(\mu) &= \lim_{\mu \rightarrow \infty} \int u(P) dF_{\mu}(P) + \beta \int_w W(P, 1) dF_{\mu}(P) \\ &> u(w) + \beta W_0. \end{aligned}$$

Hence, since $V_0(\mu)$ is strictly increasing in μ and continuous by the Mean Value Theorem, there exists a unique scalar μ_H that satisfies Equation (2). ■

A.5. Proof of Proposition 1. Guess

$$W(S, N) = \omega(\pi(S, N) + w) = \omega(N(S - w) + w).$$

Since

$$\begin{aligned} W(S, N) &= \gamma(\lambda, N) \int^{\mu_H} \max \left[\begin{array}{c} \int [\pi(S, N) + w + P - Z - w] dF_{\mu}(P) \\ + \beta \int \max \left[W \left(\frac{NS + P}{N + 1}, N + 1 \right), W(S, N) \right] dF_{\mu}(P), \\ \pi(S, N) - Z + w + \beta W(S, N) \end{array} \right] dH(\mu) \\ &\quad + (1 - \gamma(\lambda, N)) H(\mu_H) [\pi(S, N) + w + \beta W(S, N)], \end{aligned}$$

then

$$\begin{aligned} \omega(N(S-w) + w) &= (1 + \beta\omega)(N(S-w) + w) - \gamma(\lambda, N) H(\mu_H) Z \\ &\quad + \gamma(\lambda, N) \int^{\mu_H} \max\left[\mu - w + \beta\omega \int \max[P - w, 0] dF_\mu(P), 0\right] dH(\mu). \end{aligned}$$

Let

$$\omega = 1 - \beta\omega$$

or

$$\omega = \frac{1}{1 - \beta};$$

then

$$\begin{aligned} Z &= \frac{1}{H(\mu_H)} \int^{\mu_H} \max\left[\mu - w + \frac{\beta}{1 - \beta} \int \max[P - w, 0] dF_\mu(P), 0\right] dH(\mu), \\ &= \frac{1}{H(\mu_H)} \int_{\mu_L}^{\mu_H} \max\left[\mu - w + \frac{\beta}{1 - \beta} \int_{P_L} [P - w] dF_\mu(P), 0\right] dH(\mu). \end{aligned}$$

Note that this is exactly the condition that makes

$$W(S, N) = \frac{\pi(S, N) + w}{1 - \beta}.$$

That is, it makes the value function of entrepreneurs equal to the present value of the current project of the firm. So entrepreneurs are willing to pay up to Z for workers' ideas, and workers will get this price or will sell the idea to another entrepreneur. Competition for ideas among entrepreneurs then guarantees that the market price of ideas Z is determined by the condition above in equilibrium.

Then

$$V(\mu_L, S, N) = (\pi(S, N) - Z + w) + \beta W(S, N)$$

implies that μ_L satisfies

$$(A.1) \quad \int (P - w) dF_{\mu_L}(P) + \frac{\beta}{1 - \beta} \int \max[P - w, 0] dF_{\mu_L}(P) = 0,$$

and therefore is independent of S and N . Hence, $\mu_L < w$, since entrepreneurs can drop the idea next period (the maximum in the second term on the l.h.s.). P_L satisfies

$$\max[P - w, 0] > 0 \quad \text{for } P > P_L,$$

so $P_L = w$. Then μ_L is given by

$$\frac{1}{1 - \beta} \int_w (P - w) dF_{\mu_L}(P) = \int_0^w [w - P] dF_{\mu_L}(P).$$

Namely, the present value of the gains from implementing a project has to be equal to the cost.

Guess that

$$W_0 = \frac{w + f_0}{1 - \beta};$$

then

$$V_0(\mu) = \int PF_\mu(P) + \frac{\beta}{1 - \beta} \int \max[P, w + f_0] dF_\mu(P).$$

The continuation value of a worker currently working in a firm, W_0 , is then given by

$$\begin{aligned} \frac{w + f_0}{1 - \beta} = \lambda \int \max \left[\int PF_\mu(P) + \frac{\beta}{1 - \beta} \int \max[P, w + f_0] dF_\mu(P), w + Z + \beta \frac{w + f_0}{1 - \beta} \right] dH(\mu) \\ + (1 - \lambda) \left[w + \beta \frac{w + f_0}{1 - \beta} \right], \end{aligned}$$

so

(A.2)

$$f_0 = \lambda \int \max \left[\int (P - w) F_\mu(P) + \frac{\beta}{1 - \beta} \int \max[P - w - f_0, 0] dF_\mu(P), 0 \right] dH(\mu) + (1 - \lambda)Z,$$

which determines f_0 as a positive constant (given H assigns positive mass to μ s such that $\int (P - w) F_\mu(P) > 0$) and verifies our guess.

Then

$$V_0(\mu) = \frac{w}{1 - \beta} + \frac{\beta f_0}{1 - \beta} + \int (P - w) dF_\mu(P) + \frac{\beta}{1 - \beta} \left[\int \max[P - w - f_0, 0] dF_\mu(P) \right].$$

Since

$$V_0(\mu_H) = w + Z + \beta W_0,$$

μ_H is given by

$$\int (P - w) dF_{\mu_H}(P) + \frac{\beta}{1 - \beta} \int \max[P - w - f_0, 0] dF_{\mu_H}(P) = Z.$$

Two results are immediate from this expression. First, since $Z \geq 0$ and $f_0 > 0$, the l.h.s. of the above equation evaluated at $\mu_H = \mu_L$ is less than the r.h.s. given that μ_L is determined by Equation (A.1). Since the l.h.s. is increasing in μ , it follows that $\mu_L < \mu_H$. Second, since P_H is such that $P - w - f_0 = 0$, it follows that $P_H = w + f_0 > w$. ■

A.6. Proof of Proposition 2. To establish this proposition, observe that when $Z = 0$, f_0 given by Equation (A.2) is positive. Then, Equations (5) and (6) imply that if $\mu_H = \mu_L$, the l.h.s. of Equation (6) is negative and hence smaller than $Z = 0$. Now note that the l.h.s. of Equation (6) is strictly increasing and continuous in μ_H as is the r.h.s. However, observe that the derivative of the l.h.s. of Equation (6) with respect to μ_H is strictly larger than the derivative of Z with respect to μ_H and the difference of both derivatives is bounded away from 0, as Z is the *average* value of the l.h.s. of Equation (6) for the interval $[\mu_L, \mu_H]$. Therefore, by the Intermediate Value Theorem there exists a unique μ_H that solves Equation (6). Given μ_H we can immediately obtain a unique value of Z and f_0 . Furthermore, μ_L is uniquely determined by Equation (5) since the r.h.s. of Equation (5) is strictly increasing and continuous in μ_L and at $\mu_L = 0$ the r.h.s. is less than zero.

We still need to show that there exists a unique invariant distribution δ_N and an average probability of entrepreneurs buying an idea θ . To show that there exists a unique θ given an invariant distribution is immediate from Equation (3). Showing existence of a unique invariant distribution is more involved and we dedicate Section 5 to it. In Corollary 1 we show that in fact such a distribution exists and is unique. ■

A.7. Extension to CARA utility. We now proceed to prove the analog of Proposition 1 for the case of exponential utility. We are interested in this case because as in the linear case the value functions can be solved analytically, which allows us to obtain an expression for the price of an idea. Furthermore, in contrast with the linear case, with exponential utility agents are not risk neutral and, most important, without the market for ideas the thresholds would depend on the size of firms. That is, absent a market for ideas the selection of ideas would be firm specific. Hence, this case further illustrates the importance of the market for ideas in generating scale independent growth.

PROPOSITION 6. *If $u(c_t) = -ae^{-bc_t}$,*

- *the value of the firm is given by*

$$W(S, N) = \frac{u(\pi(S, N) + w)}{1 - \beta},$$

- *the value of a worker is given by*

$$W_0 = u(w)f_0$$

for some positive constant f_0 ,

- *$\mu_L(S, N)$ is independent of S and N ,*
- *the thresholds for using a project are given by $P_L(S, N) = w$ and a constant $P_H(S, N) > w$,*
- *the market price of ideas is given by*

$$Z = \frac{1}{b} \log \left[\frac{1 + \frac{\beta}{1 - \beta} \int_{\mu_L}^{\mu_H} \int_{P_L} (1 - e^{-b(P-w)}) dF_{\mu}(P) dH(\mu)}{1 - \int_{\mu_L}^{\mu_H} \int (1 - e^{-b(P-w)}) dF_{\mu}(P) dH(\mu)} \right] > 0,$$

- *and $\mu_L < \mu_H$, so some ideas are implemented within existing firms and some through spinoffs.*

PROOF. Guess that

$$W(S, N) = -ae^{-b[\pi(S, N) + w]} f(N)$$

for some function $f(N)$ independent of S . Substitute the guess to get

$$\begin{aligned} -ae^{-b[\pi(S, N) + w]} f(N) &= \gamma(\lambda, N) \int^{\mu_H} \max \left[\begin{aligned} &V(\mu, S, N) + ae^{-b[\pi(S, N) + w]} e^{bZ} \\ &+ \beta ae^{-b[\pi(S, N) + w]} f(N), 0 \end{aligned} \right] dH(\mu) \\ &+ (1 - \gamma(\lambda, N) H(\mu_H)) [ae^{-b[\pi(S, N) + w]} + \beta ae^{-b[\pi(S, N) + w]} f(N)] \\ &- \gamma(\lambda, N) H(\mu_H) [ae^{-b[\pi(S, N) + w]} e^{bZ} + \beta ae^{-b[\pi(S, N) + w]} f(N)] \end{aligned}$$

and

$$V(\mu, S, N) = \int -ae^{-b[\pi(S, N)+w]}e^{bZ}e^{-b(P-w)}dF_{\mu}(P) \\ + \beta \int \max[-ae^{-b[\pi(S, N)+w]}e^{-b(P-w)}f(N+1), -ae^{-b[\pi(S, N)+w]}f(N)]dF_{\mu}(P).$$

Then,

$$f(N) = -\frac{\gamma(\lambda, N)}{1-\beta} \int^{\mu_H} \max \left[\begin{array}{c} -\int e^{bZ}e^{-b(P-w)}dF_{\mu}(P) \\ +\beta \int \max[-e^{-b(P-w)}f(N+1), -f(N)]dF_{\mu}(P) \\ +e^{bZ} + \beta f(N), 0 \end{array} \right] dH(\mu) \\ + \frac{1}{1-\beta} + \frac{\gamma(\lambda, N)H(\mu_H)}{1-\beta}(e^{bZ} - 1),$$

which does not depend on S . This verifies our guess. Now guess that

$$f(N) = \frac{1}{1-\beta}.$$

Substituting this guess, we obtain that

$$e^{bZ} - 1 = e^{bZ} \int_{\mu_L}^{\mu_H} \int (1 - e^{-b(P-w)})dF_{\mu}(P)dH(\mu) \\ + \frac{\beta}{1-\beta} \int_{\mu_L}^{\mu_H} \int_{P_L} (1 - e^{-b(P-w)})dF_{\mu}(P)dH(\mu).$$

Hence, Z needs to satisfy this equation for our guess to be correct. Or,

$$e^{bZ} = \frac{1 + \frac{\beta}{1-\beta} \int_{\mu_L}^{\mu_H} \int_{P_L} (1 - e^{-b(P-w)})dF_{\mu}(P)dH(\mu)}{1 - \int_{\mu_L}^{\mu_H} \int (1 - e^{-b(P-w)})dF_{\mu}(P)dH(\mu)},$$

which implies that

$$Z = \frac{1}{b} \log \left[\frac{1 + \frac{\beta}{1-\beta} \int_{\mu_L}^{\mu_H} \int_{P_L} (1 - e^{-b(P-w)})dF_{\mu}(P)dH(\mu)}{1 - \int_{\mu_L}^{\mu_H} \int (1 - e^{-b(P-w)})dF_{\mu}(P)dH(\mu)} \right].$$

Note that this is exactly the condition that makes

$$W(S, N) = \frac{-ae^{-b[\pi(S, N)+w]}}{1-\beta} = \frac{u(\pi(S, N) + w)}{1-\beta},$$

so entrepreneurs are willing to pay up to this amount for workers' ideas, and workers will get this price or will sell the idea to another entrepreneur. Competition for ideas among entrepreneurs then guarantees that the condition is satisfied.

The threshold $\mu_L(S, N)$ is implicitly defined by

$$V(\mu_L(S, N), S, N) = u(\pi(S, N) - Z + w) + \beta W(S, N),$$

which can be written as

$$\int -e^{-b(P-w)} dF_{\mu_L}(P) + e^{-bZ} \frac{\beta}{1-\beta} \int \max[1 - e^{-b(P-w)}, 0] dF_{\mu_L}(P) = -1$$

or

$$\int -e^{-b(P-w)} dF_{\mu_L}(P) + e^{-bZ} \frac{\beta}{1-\beta} \int_{P_L} [1 - e^{-b(P-w)}] dF_{\mu_L}(P) = -1,$$

which determines μ_L as a constant.

P_L is then determined by

$$W\left(\frac{NS + P_L}{N+1}, N+1\right) = W(S, N)$$

or

$$\frac{e^{-b[P_L-w]}}{1-\beta} = \frac{1}{1-\beta},$$

and so $P_L = w$.

We still need to show that $P_H = w$. As before guess that

$$W_0 = -ae^{-bw}f_0,$$

$$W(P, 1) = \frac{-ae^{-bP}}{1-\beta};$$

then

$$V_0(\mu) = \int -ae^{-bw}e^{-b(P-w)} dF_{\mu}(P) + \beta \int \max\left[-ae^{-bw}e^{-b(P-w)}\frac{1}{1-\beta}, -ae^{-bw}f_0\right] dF_{\mu}(P)$$

$$= ae^{-bw} \left[-\int e^{-b(P-w)} dF_{\mu}(P) + \beta \int \max\left[-e^{-b(P-w)}\frac{1}{1-\beta}, -f_0\right] dF_{\mu}(P) \right],$$

and from the expression for W_0 ,

$$f_0 = \frac{1}{1-\beta} - \frac{\lambda}{1-\beta} \int \max \left[\begin{array}{c} \int (1 - e^{-b(P-w)}) dF_{\mu}(P) \\ + \frac{\beta}{1-\beta} \int \max[f_0 - e^{-b(P-w)}, 0] dF_{\mu}(P) \\ -e^{-bZ} \end{array} \right] dH(\mu),$$

which determines f_0 as a positive constant, where $f_0 < 1/(1-\beta)$ and verifies our guess.

The threshold μ_H is determined by

$$V_0(\mu_H) = u(w + Z) + \beta W,$$

which implies that

$$\int -e^{-b(P-w)} dF_{\mu_H}(P) + \frac{\beta}{1-\beta} \int \max[-e^{-b(P-w)} + (1-\beta)f_0, 0] dF_{\mu_H}(P) = -e^{-bZ}.$$

The equation above implies that P_H satisfies

$$e^{-b(P_L-w)} = (1-\beta)f_0 < 1,$$

since $f_0 < 1/(1-\beta)$. Hence, we conclude that $P_L > w$.

To show that $\mu_H > \mu_L$, we need to compare the equations that determine these thresholds, namely,

$$\begin{aligned} \int e^{-b(P-w)} dF_{\mu_L}(P) &= 1 + e^{-bZ} \frac{\beta}{1-\beta} \int_{P_L} [1 - e^{-b(P-w)}] dF_{\mu_L}(P), \\ \int e^{-b(P-w)} dF_{\mu_H}(P) &= e^{-bZ} + \frac{\beta}{1-\beta} \int_{P_H} [(1-\beta)f_0 - e^{-b(P-w)}] dF_{\mu_H}(P). \end{aligned}$$

But note that

$$\int_{P_H} [(1-\beta)f_0 - e^{-b(P-w)}] dF_{\mu_H}(P) < \int_{P_L} [1 - e^{-b(P-w)}] dF_{\mu_L}(P) < 1,$$

since $(1-\beta)f_0 < 1$ and $P_H > P_L$. Hence,

$$\int e^{-b(P-w)} dF_{\mu_H}(P) < \int e^{-b(P-w)} dF_{\mu_L}(P),$$

which implies that $\mu_H > \mu_L$. ■

B. Extension of Model to Allow for Entrepreneurs to Generate Ideas. In this appendix we briefly discuss an extension of the model where we assume that when spinning off, individuals do not give up the option of having new ideas. However, without this assumption spinning off would be costless, which would eliminate the market for ideas. Therefore, in this appendix, we introduce a resource cost $f > 0$ of creating a new firm, which restores the market for ideas.

The value of an entrepreneur who himself has an idea with average return μ is

$$\begin{aligned} V^E(\mu, S, N) &= \int [u(\pi(S, N) + w + P - w)] dF_{\mu}(P) \\ &\quad + \beta \int \max \left[W \left(\frac{NS + P}{N + 1}, N + 1 \right), W(S, N) \right] dF_{\mu}(P). \end{aligned}$$

An entrepreneur who buys an idea with return μ has value $V(\mu, S, N)$ as defined in the text. Note that since the current utility of the entrepreneur is different depending on whether he pays for the idea or not, we need to define both value functions.

$W(S, N)$ now satisfies the following functional equation:

$$\begin{aligned} W(S, N) &= \lambda \int \max[V^E(\mu, S, N), u(\pi(S, N) + w) + \beta W(S, N)] dH(\mu) \\ &\quad + \gamma(\lambda, N) \int^{\mu_H} \max[V(\mu, S, N), u(\pi(S, N) - Z + w) + \beta W(S, N)] dH(\mu) \\ &\quad + (1 - \lambda - \gamma(\lambda, N) H(\mu_H)) [u(\pi(S, N) + w) + \beta W(S, N)]. \end{aligned}$$

Essentially, we just added the value of getting an idea to the continuation of being an entrepreneur. Note that entrepreneurs will implement all their good ideas themselves, and so the first integral above is not truncated at μ_H . The lower threshold will be different for an idea that an entrepreneur himself gets than for ideas that he buys in the market. Let $\mu_L^E(S, N)$ be the worst of his own ideas that an entrepreneur will test/implement. This threshold is given by

$$V^E(\mu_L^E, S, N) = u(\pi(S, N) + w) + \beta W(S, N).$$

Note that in the linear case, $\mu_L^E(S, N) = \mu_L(S, N)$ because the Z cancels out from Equation (1), which determines $\mu_L(S, N)$. Whether an idea is bought or thought, the threshold $P_L = w$ at which a project is implemented remains the same because post-implementation the price paid for the idea is sunk.

The expected utility of a worker with an idea μ that decides to spin off is now given by

$$V_0(\mu) = \int u(P - f) dF_\mu(P) + \beta \int \max[W(P, 1), W_0] dF_\mu(P).$$

The rest of the worker's problem is exactly as stated in the text. Of course, now the threshold μ_H will depend on the value of the exogenous cost f . The equilibrium condition (3) now becomes

$$\lambda \sum_{N=1}^{\infty} N \delta_N = \theta \sum_{N=1}^{\infty} \tilde{\gamma}(\lambda, N) \delta_N,$$

and so for $\tilde{\gamma}(\lambda, N) = \lambda N$, $\theta = 1$.

We then have the following analog of Proposition 1:

PROPOSITION 7. *If $u(c_t) = c_t$, then in equilibrium*

- $\mu_L(S, N) = \mu_L^E(S, N)$, both are independent of S and N , and $\mu_L(S, N) < w$,
- the thresholds for using a project are given by $P_L(S, N) = P_H = w$,
- $\mu_L < \mu_H$, so some ideas are implemented within existing firms and some through spinoffs,
- and the market price of ideas is given by

$$Z = \frac{1}{H(\mu_H)} \int_{\mu_L}^{\mu_H} [\mu - w + \frac{\beta}{1 - \beta} \int \max[P - w, 0] dF_\mu(P)] dH(\mu) > 0.$$

PROOF. The proof is much simpler than in the text because now the cost of spinning off is exogenously given and does not have to be solved for. The proof that $P_L(S, N) = P_H = w$ simply follows from the fact that since the cost of spinning off is now exogenously given by f , once this cost is paid there is no reason to not carry on with a project that pays at least w . That is, there is no option value to going back to being a worker. For the rest, note that the threshold that determines μ_L is given by

$$\int (P - w) dF_{\mu_L}(P) + \frac{\beta}{1 - \beta} \int_w (P - w) dF_{\mu_L}(P) = 0,$$

independently of whether ideas are bought or thought. The threshold μ_H is now implicitly determined by

$$\int (P - w - f) dF_{\mu_H}(P) + \frac{\beta}{1 - \beta} \int_w (P - w) dF_{\mu_H}(P) = Z,$$

where, as before,

$$Z = \frac{1}{H(\mu_H)} \int_{\mu_L}^{\mu_H} \left[\mu - w + \frac{\beta}{1 - \beta} \int \max[P - w, 0] dF_\mu(P) \right] dH(\mu) > 0.$$

As discussed in the text, this system has a unique solution where $\mu_L < \mu_H$ and $Z > 0$. Note that the system is identical except for the substitution of the exogenous value f for the endogenous value f_0 . ■

Observe that unlike the model in the text, all projects that yield at least w are implemented. Therefore, there is no reason to assume that the output of projects is specific to the founder of the company. That is, firms can be sold in the market. Of course, in equilibrium there are no gains from such sales.

Finally, we need to modify the expression for transition probabilities. Now, the probability of a firm transitioning from N to $N + 1$ workers is given by

$$p(N, N + 1) = \lambda \int_{\mu_L} (1 - F_\mu(w)) dH(\mu) + \theta \lambda N \int_{\mu_L}^{\mu_H} (1 - F_\mu(w)) dH(\mu)$$

Note that in this environment, the definition of λ_H remains unchanged but that of λ_L changes accordingly. Now, the expected number of new workers in old firms is given by

$$\begin{aligned} E_t \sum_{N=1}^{L_t} \left[\int_{\mu_L} (1 - F_\mu(w)) dH(\mu) + \theta N \int_{\mu_L}^{\mu_H} (1 - F_\mu(w)) dH(\mu) \right] \lambda \delta_N \\ \equiv (\lambda_L + o(N^{-1})) L_t, \end{aligned}$$

where $o(N^{-1})$ converges to 0 as N goes to infinity. Relative to the calculations in the text, we are adding, in expectations, a constant amount of workers due to the fact that the one entrepreneur in each firm is now contributing ideas also and those ideas have a higher probability of being implemented than ideas bought in the market. This implies that growth is not scale independent: Small firms grow faster than large firms, although this dependence becomes vanishingly small as N becomes large. Therefore, the upper tail of the invariant distribution will again satisfy the statement in Proposition 5 in the text.

REFERENCES

- ANTON, J. J., AND D. YAO, "Expropriation and Inventions: Appropriable Rents in the Absence of Property Rights," *American Economic Review* 84 (1994), 190–209.
- , AND —, "Start-ups, Spin-offs, and Internal Projects," *Journal of Law, Economics and Organization* 11 (1995), 362–78.
- , AND —, "The Sale of Ideas: Strategic Disclosure, Property Rights, and Contracting," *Review of Economic Studies* 69 (2002), 513–31.
- AXTELL, R. L., "Zipf Distribution of U.S. Firm Sizes," *Science* 293 (2001), 1818–20.
- EECKHOUT, J., "Gibrat's Law for (All) Cities," *American Economic Review* 94 (2004), 1429–51.
- ERICSON, R., AND A. PAKES, "Markov-Perfect Industry Dynamics: A Framework for Empirical Work," *Review of Economic Studies* 62 (1995), 53–82.
- GABAIX, X., "Zipf's Law for Cities: An Explanation," *Quarterly Journal of Economics* 114 (2005), 739–67.
- , "The Granular Origins of Aggregate Fluctuations," *Econometrica*, 79 (2011), 733–72.
- GILSON, R. J., "The Legal Infrastructure of High Technology Industrial Districts: Silicon Valley, Route 128, and Covenants Not to Compete," *New York University Law Review* 74 (1998), 575–629.
- FRANCO, A., AND D. FILSON, "Spin-outs: Knowledge Diffusion through Employee Mobility," *Rand Journal of Economics* 37 (2006), 841–60.
- HALL, R., AND S. WOODWARD, "The Incentives to Start New Companies: Evidence from Venture Capital," NBER Working Paper 13056, April 2007.

- HOPENHAYN, H. A., "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica* 60 (1992), 1127–50.
- JOVANOVIC, B., "Selection and the Evolution of Industry," *Econometrica* 50 (1982), 649–70.
- , AND P. ROUSSEAU, "Extensive and Intensive Investment over the Business Cycle," Working paper NYU, May 2009.
- KLEPPER, S., AND S. SLEEPER, "Entry by Spin-offs," *Management Science* 51 (2005), 1291–306.
- KLETTE T. K., AND S. KORTUM, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy* 112 (2004), 986–1018.
- LUTTMER, E. G. J., "Selection, Growth, and the Size Distribution of Firms," *Quarterly Journal of Economics* 122 (2007), 1103–44.
- , "On the Mechanics of Firm Growth," Working Paper, University of Minnesota, February 2008.
- PRUSA, T., AND J. SCHMITZ, "Can Companies Maintain Their Initial Innovative Thrust? A Study of the PC Software Industry," *Review of Economics and Statistics* 76 (1994), 523–40.
- ROSSI-HANSBERG, E., AND M. L. J. WRIGHT, "Establishment Size Dynamics in the Aggregate Economy," *American Economic Review* 97 (2007), 1639–66.
- SILVEIRA, R., AND R. WRIGHT, "Search and the Market for Ideas," *Journal of Economic Theory* 145 (2010), 1550–73.
- SIMON H., AND C. BONINI, "The Size Distribution of Business Firms," *American Economic Review* 48 (1958), 607–17.
- STOKEY, N., R. E. LUCAS, JR., AND E. C. PRESCOTT, *Recursive Methods in Economic Dynamics* (Cambridge, MA: Harvard University Press, 1989).
- SUTTON, J., "Gibrat's Legacy," *Journal of Economic Literature* 35 (1997), 40–59.