

# Marginal Ideas\*

John Horton<sup>†</sup>

Amir Sariri<sup>‡</sup>

May 2, 2026

[Click here for the latest version](#)

## Abstract

We develop a model in which startups are free to pursue any idea among a collection of ideas. This competition equalizes expected private returns across those ideas, despite differing in apparent ex ante quality. Workers can either found a startup or join a startup as an employee. Their choice equalizes expected private returns across employment and entrepreneurship. We show that crowding is excessive on every idea except the “marginal idea” because entrepreneurs do not internalize business-stealing externalities. Because the business-stealing wedge grows with crowding, ideas that look unattractive ex ante have higher social returns at the margin. The idea with the highest social return is the “marginal idea”—or the worst idea that still gets funded. Pushing the margin out—notably, through public R&D that expands the set of usable ideas—lessens congestion and improves efficiency. The model generates a sufficient statistic for social efficiency, which is the elasticity of startup success with respect to number of startup entrants. We compute this elasticity cross sectionally using data from 37,818 U.S. startups across 2,045 distinct ideas. We estimate a lower-bound of  $-0.18$  for the elasticity of per-entrant success probability with respect to the number of entrants on an idea. This estimate grows to  $-0.86$  in the sample of idea niches that have at least one exit within six years of founding, meaning that a 10% increase in entry on such ideas yields only a 1.4% increase in total successful exits.

*Keywords:* entrepreneurship, startups, congestion externality, venture capital, LLMs

*JEL Classifications:* M13, L26, L22, D62, O31

---

\*Thanks to Aaron Chatterji, Ajay Agrawal, Andrey Fradkin, Ashish Arora, Ben Golub, Bill Kerr, Bo Cowgill, Boyan Jovanovic, Chuck Easley, Deepak Hegde, FM Scherer, Glen Weyl, Hal Varian, Hanna Halaburda, Ian Cockburn, Luke Stein, Matt Backus, Megan MacGarvie, Petra Moser, Ramesh Johari, Roy Radner, Shane Greenstein, Steve Tadelis, Tal Gross, Tom Åstebro, and Joe Golden for helpful comments, suggestions, and discussions. Thanks to participants at the NBER productivity seminar, the Duke Strategy Conference and the Creativity and Innovation Research Initiative seminar. We are grateful to Expected-Parrot staff for providing technical support for our LLM method implementation, and to the Media Cloud team for accommodating our unusually high demand on their servers. Thanks to Jayson Roxas for excellent research assistance and to Michelle Lindner for superb help in preparing the manuscript. We have used AI tools in this research.

<sup>†</sup>MIT & NBER

<sup>‡</sup>Purdue University

# 1 Introduction

The notion that “ideas don’t matter” is something like common wisdom in Silicon Valley. If taken literally, it is clearly false—there are terrible business ideas that should not be pursued. Silicon Valley is nearly obsessed with ideas; many if not most successful startups exploit some recent technological advance that made their idea viable. But this “ideas don’t matter” also captures an economic insight. Patents notwithstanding, anyone can pursue an idea, so the idea itself rarely distinguishes the winners. For investors and entrepreneurs considering some idea, it is either likely bad (low probability of success but you will have it to yourself) or likely good (high probability of success but you will have serious competition). The result is a kind of equilibrium in the idea market: entry equalizes expected returns across ideas, leaving every fundable idea roughly as attractive as any other *ex ante*.

A striking example of a seemingly bad idea that turned out to be good is home-sharing as pioneered by Airbnb. When Airbnb sought seed funding, most potential investors found the idea of strangers sleeping in each other’s apartments implausible. Fred Wilson of Union Square Ventures rejected investing in Airbnb, despite insistence from Paul Graham, respected YCombinator founder who had backed the company. By Airbnb’s IPO in December 2020, a half a million dollar seed investment in the company was estimated to be worth roughly \$4.8 billion.<sup>1</sup> In contrast, consider ride-sharing as a seemingly good idea that turned out to be good. The idea of booking a ride from a phone was so immediately legible to investors and entrepreneurs that it attracted dozens of well-funded competitors around the world, most of which failed.<sup>2</sup> The ride-sharing example also illustrates a general feature of technology entrepreneurship. Because business opportunities depend on widely observable technological advances, many entrepreneurs recognize the same opportunities at the same time. Collections of startups pursuing more or less the same idea are common, and experienced founders surely consider the competition they will face.

The competition between startups and investors is, of course, privately important to them, but from a welfare standpoint, our interests differ. We are less concerned with the precise winners and losers and more concerned with the overall efficiency of the system as a whole. Aside from the resources committed, what matters for innovation and economic growth is the fraction of the “idea space” that is explored. The social interest in startups is not in inching existing markets slightly closer to perfect competition, but in bringing entirely

---

<sup>1</sup>See the original email exchange at <https://paulgraham.com/airbnb.html> and Sequoia Capital’s blog at <https://bit.ly/4aWMZSv>.

<sup>2</sup>Sidecar and Hailo, which raised \$35 million and \$100 million respectively, shut down in 2015 and 2016. Gett spent \$200 million to acquire Juno in 2017 and closed the service two years later. See <https://bit.ly/30ldLv6>, <https://bit.ly/4tDKxrz>, and <https://bit.ly/4qIYYrG>.

new products and services to market that are newly possible. With this framing, our interest is less in the “best” ideas that entrepreneurial self-interest promises to cover, than in the “marginal” idea, or the worst idea that still gets funded.

The most important research questions are those about the entire system, not just the individual startups. What determines this marginal idea, in equilibrium? What determines how much entrepreneurial entry a given idea attracts, and what are the consequences of this entry for the startups involved? When do we have too much or too little entrepreneurship, at least within the world implied by the model? What policy levers would push the system towards the socially optimal allocation? We address these questions by developing a tractable model of a high technology entrepreneurship cluster.

Startups require venture capital, and entrepreneurs pursuing the same business idea compete in winner-take-all product markets. “Engineers” choose between founding a startup and working as employees at those same ventures. For those that pursue entrepreneurship, better ideas attract proportionally more entrants until per-firm success probabilities equalize across all pursued ideas. This equalization creates a business-stealing externality that individual entrepreneurs do not internalize. Because the product markets are winner-take-all, free entry into technology entrepreneurship generates a congestion externality. We show that the welfare loss is characterized by a single sufficient statistic, the elasticity of startup success probability with respect to the number of entrants. A critical object is the seemingly worst idea that still gets funded, what we call the marginal idea.

The model yields four main results. First, the decentralized equilibrium generates excessive entry whenever a particular elasticity is negative. This elasticity is the probability of success of a startup with respect to the number of entrants pursuing the same idea, which can be measured. Excess entry occurs because entrepreneurs capture their full private expected return but ignore the crowding externality they impose on others pursuing the same idea. Second, when startups face idiosyncratic execution risk that justifies multiple attempts per idea, the excess entry concentrates disproportionately on the highest-quality ideas. The socially optimal allocation, by contrast, is concave in idea quality, spreading more effort toward marginal ideas. The divergence between decentralized and optimal allocations is steepest at the top of the quality distribution, creating what we call a “best-idea trap.” Third, we show that engineer wages are a sufficient statistic for total system output, providing a directly observable metric for evaluating the health of an entrepreneurial ecosystem without requiring knowledge of the underlying idea distribution. Finally, the model shows that the inefficiency can be lessened by public R&D that expands the set of usable ideas. It is not simply that more ideas means more chances of entirely new products—it is that more ideas also lessens congestion. Expanding the supply of economically distinct ideas directly reduces

congestion by shifting entry from the intensive margin to the extensive margin.

We complement the theory with descriptive results from two distinct samples. The descriptive empirics speak to the entry and congestion predictions. Using data on 37,818 US technology startups and a new LLM-based measurement approach that assigns startups to 2,045 product-market niches, we track entry and exit patterns across the idea quality distribution. On entry, we find that higher perceived idea quality, measured by media attention, is associated with greater entry, yet per-firm success probabilities are statistically indistinguishable from flat across the quality distribution. The estimated elasticity of success with respect to entry is  $-0.18$  in the full sample of 2,045 niches, rising in magnitude to  $-0.86$  on the 458 niches with at least one exit within six years of founding. We treat the full-sample estimate as a conservative lower bound on the true congestion elasticity and the restricted-sample estimate as an upper bound. Both fall within the theoretical  $[-1, 0]$  continuum the model identifies, between the novel-idea benchmark ( $\eta = 0$ , where every entrant finds a previously unexplored opportunity) and the complete-crowding benchmark ( $\eta = -1$ , where every entrant merely displaces an existing competitor).

We also find that at the idea level, the probability that a niche produces at least one successful exit is concave in entry, consistent with the diminishing returns that drive the excess-entry result. Using externally timed AWS cost and tooling shocks, we also show that cloud-exposed niches experience more entry after the 2014 cloud wave, without corresponding gains in downstream exit outcomes. In a complementary sample of Y Combinator startups, we find that aggregate batch size does not predict startup failure, but the number of thematically similar companies funded in the same batch does. This distinction corroborates the model’s central mechanism: congestion is generated by dense entry into a specific opportunity, not by entrepreneurial activity in the aggregate.

A central challenge for testing congestion in entrepreneurship has been the absence of product-market niche definitions at the startup level. Published methods for measuring competitive proximity rely on SEC filings such as 10-K product descriptions available only for public firms (Hoberg & Phillips, 2010). Building on recent advances in the use of LLMs in social science research (Manning *et al.*, 2024), we address this gap with a new LLM-based approach that infers the competitive neighborhood of each startup from its brief business description. The method asks an LLM agent, embedded in the role of a Google AdWords specialist, to generate the set of customer search intents each firm would pay to reach, then clusters firms by the overlap in these keyword portfolios. This approach enables the first measurement of the congestion elasticity  $\eta$  at the product-market niche level across the population of US technology startups, placing the ecosystem on the  $[-1, 0]$  continuum that the theory identifies.

The remainder of the paper proceeds as follows. [Section 2](#) reviews related work. [Section 3](#) presents the model. [Section 4](#) derives the equilibrium and the congestion elasticity  $\eta$ . [Section 5](#) shows that the decentralized allocation features excess entry concentrated on the best ideas. [Section 6](#) evaluates the effect of an AI productivity shock on equilibrium wages, profits, probability of success, and the share of labor in entrepreneurship. [Section 7](#) introduces the LLM-based measurement of product-market niches and provides empirical results. We conclude with directions for future research in [Section 8](#).

## 2 Free Entry and Occupational Choice

Venture-backed, growth-oriented entrepreneurship is a small fraction of all entrepreneurship ([Åstebro \*et al.\*, 2014](#)), and much of what we know about the conventional variety does not apply to it. Conventional entrepreneurs enter established markets, often self-financed, with entry positively correlated with wealth ([Evans & Jovanovic, 1989](#); [Blanchflower & Oswald, 1998](#); [Holtz-Eakin \*et al.\*, 1994](#)) and downside risk managed through firm size, capital structure, and the option to default ([Herranz \*et al.\*, 2013](#)).<sup>3</sup> Innovative entrepreneurship looks different. Founders confront extreme payoff skewness and depend on specialized equity finance ([Lerner & Nanda, 2020](#)). VC syndication and staged finance decouple entry from founder wealth and make competitive dynamics among simultaneously entering startups central to returns.

Our model shares commonalities with the free-entry framework of [Mankiw & Whinston \(1986\)](#), where entry is socially excessive because entrants steal business from incumbents. In our setting, the “business” being stolen is the probability of being the winning startup in a product market. Schumpeterian growth models formalize similar business-stealing externalities in the context of innovation ([Aghion & Howitt, 1992](#); [Akcigit & Kerr, 2010](#); [Acemoglu \*et al.\*, 2013](#)), and multi-product firm models feature entry on specific product lines ([Klette & Kortum, 2004](#)). In those models, however, firms arrive on product lines by chance rather than choosing among visible opportunities.

A growing literature studies how entrepreneurs resolve uncertainty through costly experimentation ([Agrawal \*et al.\*, 2021, 2026](#)) and how external advisors shape the learning process. [Sariri \(2025\)](#) shows that VC and angel mentoring improves startup market performance by shifting early-stage founders’ task priorities from execution toward testing idea viability. These studies focus on within-idea uncertainty about how to execute, whereas our model focuses on across-idea uncertainty about where to enter.

Lastly, entrepreneurial strategy research emphasizes multiple commercialization paths

---

<sup>3</sup>[Hurst & Lusardi \(2004\)](#) question the liquidity-constraint interpretation of the wealth-entry correlation.

and commitment under uncertainty (Gans *et al.*, 2019), complemented by the entrepreneurial finance work examining how experimentation and capital availability determine the ideas that get pursued (Ewens *et al.*, 2018; Lerner & Nanda, 2020; Howell, 2020). We extend these studies by recognizing the congestion externality that arises when multiple entrepreneurs pursue the same opportunity.

### 3 Model Setup: Engineers, Ideas, and Venture Capital

The theory builds on two lines of work. Occupational choice models in the tradition of Kihlstrom & Laffont (1979), Lucas (1978), and Lazear (2004) ask who becomes an entrepreneur but treat the opportunity as given and focus on the margin between employment and self-employment. Free-entry models following Mankiw & Whinston (1986) show that entry can be socially excessive when entrants steal business from incumbents, and Schumpeterian growth models feature similar dynamics on product lines (Aghion & Howitt, 1992; Klette & Kortum, 2004), but firms in those models arrive on product lines by chance rather than choosing among visible opportunities. Our model combines the occupational choice margin with endogenous idea selection, and the interaction between these two margins generates the congestion externality that neither tradition captures alone. Recent experimental evidence from Bryan *et al.* (2026) reinforces the mechanism, showing that workers in startup labor markets respond strongly to common quality signals.

Most occupational models of entrepreneurship posit some individual characteristic that determines selection, whether managerial skill (Lucas, 1978), balanced skills (Lazear, 2004), risk tolerance (Kihlstrom & Laffont, 1979), or opportunity spotting (Holmes & Schmitz Jr, 1990).<sup>4</sup> For innovative entrepreneurship, these characteristics matter less. Risk aversion is central in Kihlstrom and Laffont, but VC equity financing eliminates the personal debt obligations and wealth constraints that dominate conventional entrepreneurship (Evans & Jovanovic, 1989), so the primary cost of founding is the multi-year opportunity cost of the founder’s time. Managerial ability determines selection in Lucas, but technology founders must first be technically excellent, and successful startups bring in professional management as they scale.<sup>5</sup> We therefore treat founders as *ex ante* symmetric and focus on how their entry decisions interact in equilibrium.

---

<sup>4</sup>These models share a Knightian and Roylian logic (Knight, 1921; Roy, 1951) in which financial returns, mediated by individual characteristics, drive the occupational margin. Gromb & Scharfstein (2002) also model entrepreneurs who differ in ability, focusing on the relative returns to internal versus external innovation.

<sup>5</sup>As Lazear (2004) writes, conventional innovation “may be as seemingly minor as recognizing that a particular street corner would be a good location for a dry cleaner.” Technology entrepreneurs create entirely new products, making idea selection central rather than individual characteristics.

The model captures an entrepreneurial cluster as three interconnected markets: the market for venture capital, the labor market for high-skilled individuals (“engineers”), and the product markets that successful startups serve. There is a mass  $S$  of engineers who must choose between founding a startup as an entrepreneur or joining an established startup as an employee. This occupational choice determines both the supply of entrepreneurial ventures and the availability of skilled labor needed to scale successful firms. [Table 1](#) collects the notation used throughout the paper.

A startup is founded by a single engineer implementing a single business idea. Ideas differ in their probability of product market success,  $q \in [0, 1]$ , distributed with pdf  $f(\cdot)$  over a mass  $\kappa$  of potential ideas. These success probabilities are common knowledge, capturing the notion that certain technological or market opportunities are more promising than others. Technology entrepreneurs are typically exploiting recent, publicly observable advances. This logic parallels the evidence on simultaneous discovery in science ([Merton, 1957](#); [Lemley, 2011](#); [Marx & Hsu, 2022](#)) and implies that many entrepreneurs recognize the same opportunities at the same time, making common signals the natural driver of correlated entry.

Any entrepreneur is free to pursue any idea, but multiple entrepreneurs pursuing the same idea must compete in a winner-take-all product market where at most one startup succeeds, and then only if the underlying idea proves viable. Successful startups generate revenue  $R$  through a production function  $\phi(l)R$ , where  $l$  represents the number of engineers hired as employees and  $\phi : [0, S] \rightarrow [0, 1)$  is increasing and concave with  $\phi'(l) > 0$  and  $\phi''(l) < 0$ . The assumption that  $\lim_{l \rightarrow S} \phi(l) = 1$  ensures that when nearly all engineers work as employees, the marginal return to entrepreneurship exceeds that of employment, guaranteeing an interior equilibrium. We model the long-run equilibrium of the cluster rather than adjustment dynamics. We make this choice because a static formulation buys tractability without sacrificing the entry and congestion forces that are the paper’s focus.

### 3.1 Entrepreneurs Sort Across Ideas

When  $n$  entrepreneurs pursue an idea with success probability  $q$ , each faces an individual success probability of  $q/n$ . We adopt winner-take-all as a simplifying assumption that captures non-commodity markets with substantial market power, driven by forces such as network effects, patent races, or first-mover advantages. The assumption is not critical to the model. An oligopoly generalization is more complex without being more interesting, since the congestion and concavity predictions carry through.<sup>6</sup>

---

<sup>6</sup>The assumption is also testable. [Section 7.4](#) shows the data reject the strict single-winner prediction but are consistent with a small number of winners per idea, and the qualitative congestion and concavity predictions survive this generalization.

For a given threshold  $q_0$ , the total number of funded entrepreneurs is:

$$E = \kappa \int_{q_0}^1 \frac{q}{q_0} f(q) dq \quad (1)$$

The equilibrium value of  $q_0$  emerges endogenously as engineers optimize their occupational choice, balancing the returns from entrepreneurship against wage employment. This equalization of success probabilities across pursued ideas captures a key insight. In particular, while better ideas are inherently more attractive, they draw proportionally more entrants until individual returns equalize. The marginal idea with quality  $q_0$  attracts exactly one entrepreneur. The marginal idea with quality  $q_0$  attracts “one” entrepreneur in the sense that the mass of entrants on idea  $q_0$  equals  $q_0/q_0 = 1$ , since both ideas and engineers are continua. Thus, the marginal idea is the natural boundary between pursued and unpursued opportunities.

## 4 Equilibrium Characteristics

### 4.1 Financing and Wages

Entrepreneurs require seed capital  $c$  to start their ventures, obtaining this funding from venture capitalists in exchange for equity. The seed stage is where the occupational choice and idea-selection margins that drive our results are active. We abstract from follow-on rounds, where much of the residual uncertainty about  $q$  and execution quality is resolved.

In equilibrium, VCs earn a required return  $r$  on their investment, implying that the entrepreneur’s retained equity share  $e$  satisfies:

$$(1 + r)c = (1 - e)q_0\pi \quad (2)$$

where  $\pi$  represents the profits of a successful startup. These profits equal revenue minus labor costs:  $\pi = \phi(l^*)R - wl^*$ , where  $l^*$  is the profit-maximizing number of employees and  $w$  is the equilibrium wage (see Appendix B.3 for proof that profits are always positive.).

The occupational indifference condition requires that expected returns from entrepreneurship equal the wage from employment:

$$w = eq_0\pi = q_0\pi - C \quad (3)$$

where  $C \equiv (1 + r)c$  represents the total cost of founding a startup, including the opportunity cost of capital. This condition, combined with labor market clearing where  $(1 - g)S$  engineers work as employees for  $q_0E$  successful startups determines the equilibrium fraction of entrepreneurs:

$$g = \frac{1}{1 + q_0l^*} \quad (4)$$

**Table 1:** Notation used in the theory

Symbol	Description
<i>Primitives</i>	
$S$	Mass of engineers in the cluster
$\kappa$	Mass of potential business ideas
$R$	Product-market revenue scale (attained revenue is $\phi(l)R$ )
$c$	Seed capital required to found a startup
$r$	Venture capitalist's required return on seed capital
$C \equiv (1 + r)c$	Total (amortized) cost of founding a startup
$\alpha$	Labor-augmenting productivity multiplier; baseline $\alpha = 1$ (AI shock)
$q$	Idea-specific probability of product-market success ("quality"), $q \in [0, 1]$
$f(q)$	Density of idea quality over $[0, 1]$
$\phi(l)$	Per-firm production function; increasing, strictly concave, with $\phi : [0, S] \rightarrow [0, 1)$ and $\lim_{l \rightarrow S} \phi(l) = 1$
$P(n)$	Probability at least one of $n$ startups succeeds on a viable idea; $P(n) = 1 - (1 - p)^n$
<i>Endogenous equilibrium objects</i>	
$q_0$	Equilibrium per-firm success probability (quality of the marginal pursued idea)
$E$	Mass of entrepreneurs, $E = \kappa \int_{q_0}^1 (q/q_0) f(q) dq$
$g$	Fraction of engineers in entrepreneurship, $g = E/S = 1/(1 + q_0 l^*)$
$l^*$	Employees per successful startup (profit-maximizing)
$n(q) = q/q_0$	Entrepreneurs pursuing an idea of quality $q$ in the decentralized equilibrium
$n^*(q)$	Social-planner allocation of entrepreneurs across ideas
$w$	Equilibrium engineer wage
$e$	Founder's retained equity share; pinned by $(1 + r)c = (1 - e)q_0\pi$
$\pi$	Profit of a successful startup, $\pi = \phi(l^*)R - wl^*$
<i>Elasticities and ratios</i>	
$\eta_E^{q_0}$	Elasticity of success probability $q_0$ w.r.t. mass of entrepreneurs (congestion elasticity); $\in [-1, 0]$
$ \eta_w^D $	Magnitude of the wage elasticity of firm-level labor demand
$h \equiv \kappa q_0 f(q_0)/E$	Density of ideas at the margin relative to the mass of entrepreneurs

Notes: Throughout the paper,  $\eta_Y^X$  denotes the elasticity of variable  $X$  with respect to  $Y$ , with  $|\cdot|$  indicating magnitude.

The model thus endogenously determines wages, equity splits, and the allocation of engineers between occupations. Higher startup costs  $C$  reduce entrepreneurial returns, shifting engineers toward employment. Conversely, greater product market opportunities  $R$  increase both wages and entrepreneurial equity, with the net effect on occupational choice depending on the elasticity of labor demand.

## 4.2 Existence and Uniqueness of Equilibrium

The model's equilibrium requires that engineers be indifferent between occupations, that VCs earn their required return, and that all markets clear. The returns to each occupation depend on the fraction choosing entrepreneurship: as  $g$  increases, more startups compete for ideas (lowering  $q_0$ ) while fewer engineers supply labor (raising  $w$ ). These opposing forces ensure a unique crossing point.

**Proposition 1** *With a sufficiently small ratio of startup costs to potential revenue,  $C/R$ , a unique equilibrium exists.*

*Proof.* See Appendix B.1.

The proof establishes that as  $g \rightarrow 0$ , the returns to entrepreneurship exceed those from employment since there is minimal idea competition and  $q_0$  approaches the quality of the best available idea. Conversely, as  $g \rightarrow 1$ , employment dominates since excessive competition drives success probabilities toward zero while labor scarcity pushes wages arbitrarily high. The monotonicity of returns in  $g$  ensures these curves cross exactly once in the unit interval, pinning down the unique equilibrium allocation.

## 4.3 When Does New Entry Expand the Frontier?

A critical parameter governing the model's behavior is the elasticity of startup success probability with respect to the number of entrepreneurs. To derive this elasticity, consider how  $q_0$  adjusts when the number of entrepreneurs increases by  $dE > 0$ . Some of these new entrepreneurs pursue ideas that were already being pursued by incumbents, while others pursue "new" ideas with quality  $q < q_0$  that were previously below the viability threshold.

When  $dE$  entrepreneurs enter, ideas that were already pursued (those with  $q > q_0$ ) each receive  $q/q_0^2 \cdot dq_0$  additional entrepreneurs, as the threshold  $q_0$  falls by  $dq_0$ . The total increase in entrepreneurs on existing ideas is thus  $\kappa \int_{q_0}^1 (q/q_0^2) dq_0 \cdot f(q) dq = E/q_0 \cdot dq_0$ . Additionally, a mass  $\kappa f(q_0) dq_0$  of entrepreneurs pursue previously unpursued ideas in the interval  $[q'_0, q_0]$ , where  $q'_0 \equiv q_0 - dq_0$  is the new threshold. Therefore:

$$dE = \kappa f(q_0) dq_0 + \frac{E}{q_0} dq_0 \tag{5}$$

After rearranging, we obtain the elasticity of startup success probability with respect to entrepreneurs:

$$\eta_E^{q_0} = -\frac{1}{1 + \kappa f(q_0)q_0/E} \quad (6)$$

This elasticity captures the extent of idea competition. When  $\eta_E^{q_0} = -1$  (which occurs when  $\kappa f(q_0) \rightarrow 0$ ), there are no unpursued ideas at the frontier, and additional entrepreneurs merely crowd existing ideas without exploring new opportunities. Entrepreneurship becomes a zero-sum game. Conversely, when  $\eta_E^{q_0} = 0$  (which occurs when  $\kappa f(q_0) \rightarrow \infty$ ), each marginal entrepreneur pursues a previously unexplored idea, eliminating business-stealing effects.

The parameter  $\kappa f(q_0)$  represents the density of ideas at the margin. Industries with high startup costs naturally have fewer entrepreneurs and thus more unexplored ideas at the margin, leading to less elastic success probabilities. In contrast, sectors with low entry barriers (e.g., software) exhibit highly elastic success probabilities because low barriers attract a large mass of engineers into entrepreneurship, leaving few unexplored ideas at the frontier relative to the number of entrants.

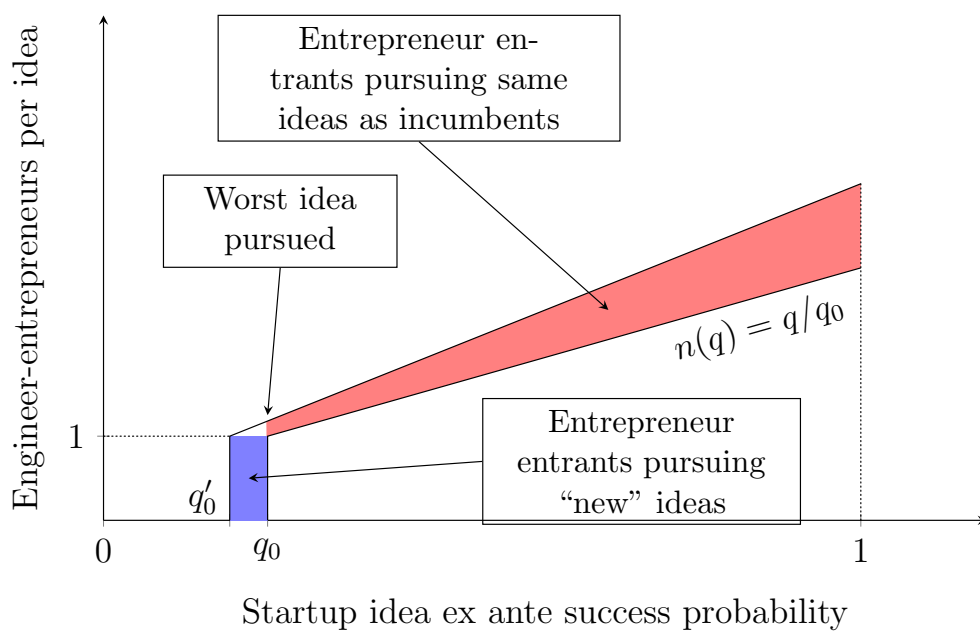
#### 4.4 What Startups Exist in Equilibrium?

**Figure 1** illustrates how entrepreneurs distribute across ideas of varying quality. The upward-sloping line  $n(q) = q/q_0$  shows the number of entrepreneurs per idea in equilibrium, where higher-quality ideas attract proportionally more entrants. The marginal idea at  $q_0$  attracts exactly one entrepreneur, while an idea with success probability  $q = 0.8$  would attract  $0.8/q_0$  entrepreneurs if, for instance,  $q_0 = 0.2$  (implying four entrepreneurs on that idea).

Let  $dq_0 > 0$  denote the magnitude of the decrease in  $q_0$  (so that the new threshold is  $q_0 - dq_0$ ). Following an increase  $dE$  in the number of entrepreneurs, two effects occur. First, existing ideas attract additional entrants with each idea of quality  $q$  receiving  $q/q_0^2 \cdot dq_0$  more entrepreneurs as competition intensifies. Second, previously unpursued ideas below the original threshold become viable as  $q_0$  shifts leftward by  $dq_0$ . The shaded regions in **Figure 1** decompose these two effects. The area above the original allocation line represents crowding on existing ideas, while the area to the left of the original  $q_0$  represents entry on newly viable ideas.

The relative importance of these margins determines the elasticity  $\eta_E^{q_0}$  and thus the efficiency properties of the equilibrium. When most new entrepreneurs pursue previously unexplored ideas (the extensive margin dominates), the system efficiently expands the range of experimentation. When they primarily crowd existing opportunities (the intensive margin dominates), entry becomes socially wasteful as entrepreneurs merely redistribute rather than create success probabilities. Formally,  $|\eta_E^{q_0}|$  equals the share of marginal entry absorbed by

**Figure 1:** The allocation of entrepreneurs over business ideas of varying quality, before and after a positive shock of  $dE$  new entrepreneurs



Notes: This figure illustrates the number of entrepreneurs per idea before and after an increase of  $dE$  new entrepreneurs. Pre-shock, each idea with quality greater than  $q_0$  received  $q/q_0$  entrants, where  $q_0$  was the worst idea still funded. After entry of additional entrepreneurs, there are both (1) more entrepreneurs per idea and (2) more entrepreneurs pursuing previously un-pursued ideas, pushing down  $q_0$ .

the intensive margin.

## 4.5 Engineer Welfare Is Social Welfare

The entrepreneurial system has  $q_0 E$  successful startups, each generating revenue  $R\phi(l^*)$ . The social cost of generating this revenue is the total startup cost,  $EC$ . The net output of the system, total revenue minus total founding costs, equals the total number of engineers times the market wage:

$$Sw = q_0 ER\phi(l^*) - EC. \quad (7)$$

**Proposition 2** *The net social benefit of the entrepreneurial system is  $Sw$ .*

*Proof.* See Appendix B.12.

The proof follows from the occupational indifference condition and market clearing. Because every engineer earns  $w$  regardless of occupation, and because VC returns are pinned at  $r$ , the wage absorbs all variation in the system's primitives (idea supply, startup costs, and product market size). A policymaker who observes only the engineer wage can rank alternative states of the economy without knowledge of the underlying idea distribution, the number of entrepreneurs, or the success probability. The comparative statics in Appendix A confirm this role. A positive shock to the supply of ideas raises wages (Proposition 9), while a positive shock to the supply of engineers lowers them (Proposition 8). Each shock's effect on total output can be read off the resulting change in  $w$ .

## 5 Which Ideas Are Crowded the Most?

To assess efficiency, consider the social value of moving a marginal engineer from employment to entrepreneurship. The social cost is the foregone wage  $w$ . The social benefit depends on whether the marginal entrepreneur pursues a new idea or crowds an existing one. From our earlier analysis, the probability that a marginal entrepreneur pursues a previously unexplored idea is  $1 + \eta_E^{q_0} = \kappa f(q_0) dq_0 / dE$ . If this entrepreneur successfully develops a new product (probability  $q_0$ ), society gains revenue  $\phi(l^*)R$  but incurs labor costs  $wl^*$ , yielding net benefit  $\pi = \phi(l^*)R - wl^*$ . The social cost of the startup attempt is  $C$  regardless of success.

The social planner would be indifferent about the marginal engineer's occupation when:

$$(1 + \eta_E^{q_0})(q_0\pi) - C = w \quad (8)$$

However, the private indifference condition from the decentralized equilibrium is:

$$q_0\pi - C = w \quad (9)$$

These conditions coincide only when  $\eta_E^{q_0} = 0$ , leading to our key efficiency result:

**Proposition 3** *The decentralized allocation of engineers to entrepreneurship and employment is efficient only when the startup success probability is completely inelastic with respect to the number of entrepreneurs ( $\eta_E^{q_0} = 0$ ).*

*Proof.* See Appendix B.13.

Entering entrepreneurs capture the full private return  $eq_0\pi$  but ignore the negative externality ( $|\eta_E^{q_0}|)(q_0\pi)$  imposed on existing entrepreneurs pursuing the same ideas. This business-stealing effect, combined with the fixed cost  $C$ , leads to excessive entry whenever  $|\eta_E^{q_0}| > 0$ . The magnitude of inefficiency is directly tied to the elasticity. Highly elastic success probabilities indicate severe overcrowding, while inelastic probabilities suggest the margin of new idea exploration remains productive.

## 5.1 Excess Entry on the Best Ideas

The inefficiency identified in Proposition 3 assumes that one startup suffices to realize an idea's full social value. In practice, startups fail for idiosyncratic reasons such as execution errors, timing, or team-specific factors, suggesting potential value in multiple attempts per idea. To examine this, let  $P(n)$  denote the probability that at least one of  $n$  startups succeeds conditional on idea viability. Under the natural assumption that individual failures are independent with success probability  $p$ , we have  $P(n) \approx 1 - \exp(-pn)$ .

The social planner's problem is to allocate  $E$  entrepreneurs across ideas to maximize expected successful products:

$$\max_{n(\cdot)} \kappa \int_{q_0}^1 qP(n(q))f(q)dq \quad \text{subject to} \quad \int_{q_0}^1 n(q)f(q)dq = E \quad \text{and} \quad n(q_0) = 1 \quad (10)$$

The first-order condition yields  $qP'(n^*(q)) = k_0$  for some constant  $k_0$ , implying that the marginal return to an additional entrepreneur should equalize across all ideas. With  $P(n) \approx 1 - \exp(-pn)$ , this gives:

$$n^*(q) = k_1 - \frac{1}{p} \log \left( \frac{k_1}{q} \right) \quad (11)$$

where  $k_0$  and  $k_1$  are constants determined by the constraint and boundary condition.

This optimal allocation is concave in  $q$ :  $n^{*''}(q) = -1/(pq^2) < 0$ . In contrast, the decentralized allocation  $n(q) = q/q_0$  is linear in  $q$ .

**Proposition 4** *If individual startup success probabilities are independent and identically distributed, the decentralized equilibrium exhibits excess entry on high-quality ideas relative to the social optimum.*

We work through the argument here because it exposes the mechanism most directly. Since both allocations must satisfy the same constraint and boundary condition, and since  $n^*(q_0) > n'(q_0) = 1/q_0$  (the optimal allocation is steeper at the margin), the curves must cross at some  $\hat{q} > q_0$ . For all ideas with quality  $q > \hat{q}$ , the decentralized equilibrium generates excessive entry  $n(q) > n^*(q)$ . The intuition is that entrepreneurs consider only their private returns  $eq\pi/n(q)$ , not the crowding externality imposed on other entrants. This externality is most severe for the highest-quality ideas that attract the most entrepreneurs.

This result provides theoretical support for contrarian investment strategies in a specific sense: high- $q$  ideas suffer the greatest excess entry, while marginal ideas near  $q_0$  face less competition. A patient investor who can systematically identify marginal ideas therefore operates in a thinner competitive field—though formalizing this as a risk-adjusted return advantage would require extending the model to heterogeneous investors.

## 6 How Does an AI Productivity Shock Change the Equilibrium?

The rise of AI coding assistants is a salient productivity shock for the population our model centers on. Software engineering ranks among the occupations most exposed to generative AI (Eloundou *et al.*, 2024). Our framework offers a direct lens on how such a shock reshapes the cluster. We introduce it as a labor-augmenting productivity multiplier  $\alpha \geq 1$  on the employee production function,  $\phi(l) \rightarrow \phi(\alpha l)$ , with  $\alpha = 1$  the baseline. One engineer now produces as much output as  $\alpha$  engineers did before.

Agrawal *et al.* (2019) distinguish prediction tasks, where AI substitutes for labor, from judgment tasks, where it complements labor. In our setting the natural mapping is that engineers produce prediction-heavy output while founders supply the judgment that selects which problem to pursue. Labor-augmenting productivity at engineer tasks therefore captures the first-order effect of contemporary coding AI agents. The firm’s problem is  $\max_l \phi(\alpha l)R - wl$ , with first-order condition  $\alpha\phi'(\alpha l^*)R = w$ . Defining efficiency units of labor  $u^* \equiv \alpha l^*$ , the optimum depends only on the efficiency wage  $w/\alpha$ :  $u^* = \psi(w/\alpha)$ , where  $\psi \equiv (\phi')^{-1}(\cdot/R)$ . Human labor per successful firm is  $l^* = \psi(w/\alpha)/\alpha$ , and profits inherit the same structure,  $\pi(w, \alpha) = \Pi(w/\alpha)$ . The shock is not literally isomorphic to an expansion of engineer supply (Proposition 8) since startup costs  $C$  are nominal and do not scale with  $\alpha$ , and the mass of potential entrepreneurs is not inflated in efficiency units because founders are human.

Let  $h \equiv \kappa q_0 f(q_0)/E$  denote the density of ideas at the margin relative to the mass of entrepreneurs; it governs how much the idea threshold moves when the entrepreneur count

changes.

**Proposition 5** *An increase in engineer productivity: (1) has an ambiguous effect on the wages of engineers, positive when  $g|\eta_w^D| + eh > g(1 - e)$ , (2) has an ambiguous effect on the retained equity of entrepreneurs, with the same sign as the wage effect, (3) has an ambiguous effect on expected profits, with the same sign as the wage effect, (4) raises realized profits, (5) has an ambiguous effect on the startup probability of success, with sign equal to that of  $|\eta_w^D|g - 1$ , and (6) has an ambiguous effect on entrepreneurship, with sign opposite to (5).*

*Proof.* See Appendix [B.11](#).

Two conditions govern the proposition. The first,  $g|\eta_w^D| + eh > g(1 - e)$ , determines whether the productivity gain passes through to engineer wages. It is satisfied whenever  $|\eta_w^D| \geq 1 - e$ , and fails only when labor demand is very inelastic, the idea margin is tight, and the VC share is large. For software, if  $|\eta_w^D|$  is of order one and the entrepreneurship share  $g$  is modest, then the condition holds and wages, retained equity, and expected profits all rise. The second condition,  $|\eta_w^D|g \leq 1$ , determines the direction of the congestion and entrepreneurship effects. It compares the per-firm labor-demand response to the entrepreneurial reallocation margin: when the wage margin responds more aggressively than the entry margin, the ecosystem absorbs less of the shock through expanded firm size and more through tighter idea competition and a higher entrepreneurship share. Our cross-sectional estimate  $\eta_E^{q_0} \approx -0.18$ , rising to  $-0.86$  on the viable-ideas subsample, places the US startup ecosystem in a congested regime, with few unpursued ideas at the frontier on the most promising opportunities; whether an AI shock amplifies or relaxes this congestion requires the separate calibration  $|\eta_w^D|g \leq 1$ . Under a modest entrepreneurship share and a labor-demand elasticity of order one, the condition  $|\eta_w^D|g < 1$  holds, and the AI shock pushes engineers toward entrepreneurship while intensifying idea competition.

**Corollary 6** *As  $\alpha \rightarrow \infty$ ,  $l^* \rightarrow 0$  and  $g \rightarrow 1$ , the economy converges to an entrepreneurship-only limit.*

The limit formalizes the intuition that sufficiently powerful tools free every engineer to found a firm. Two scope restrictions should be kept in mind. First, the model is static, so these are long-run comparative statics rather than transition dynamics. Second, the labor-augmenting specification treats AI as an employee-side tool; a complementary exercise in which AI also lowers  $C$ -augmenting founder judgment—reinforces the entrepreneurship effect without changing the sign of any result above. The broader point is that the congestion elasticity  $\eta_E^{q_0}$  governs how the ecosystem absorbs an AI shock, just as it governs responses to startup costs, supply of engineers, supply of ideas, and the extent of the product market (Propositions [7](#), [8](#), [9](#), [10](#)).

## 7 Empirical Exploration of Idea Competition

The theory yields two empirical claims: free entry equalizes per-firm success across pursued ideas, and excess entry is most severe on the most popular opportunities. Testing these predictions on US technology startups requires a firm-level measure of product-market niches. Existing methods for measuring product similarity rely on 10-K disclosures and do not extend to private firms. We close this gap with an LLM-based niche definition applied to 37,818 startups founded between 2012 and 2018 and grouped into 2,045 niches.

The model requires three empirical objects: a population of entrants, an assignment of each entrant to a product-market niche, and an observable correlate of idea quality. We construct all three from Crunchbase, a commercial repository of startup and private equity records increasingly used in academic research (Marx & Hsu, 2022; Howell, 2017), supplemented by news coverage for quality measurement.

### 7.1 Defining Entrants

Our analysis sample contains 37,818 US technology startups founded between 2012 and 2018. We arrive at this number through a sequence of filters. Since Crunchbase records are largely self-reported, the raw data include shell entries, abandoned profiles, and non-ventures. We exclude organizations that lack a homepage URL, that closed within twelve months of founding, or that report fewer than fifty characters of business description. We further drop consulting firms, rental companies, and property management businesses, whose competitive dynamics fall outside the model.

The model describes entry and competition mediated by technical labor, scalable products, and rapid imitation. A local service business operates under different economics. To align the sample with the theory, we apply an LLM classifier that reads each firm’s description and returns a binary technology-company label (see Appendix Figure D2 for prompt details). After restricting to technology firms, we cluster startups into product-market niches and drop singleton clusters with no competitive neighborhood with any other entrant.<sup>7</sup> Finally, we require that each niche contain at least one exit (IPO or acquisition) at any point in our observation window, since the model’s predictions concern markets where success is observable.

We choose six years since it is the maximum observation window available. The last

---

<sup>7</sup>A singleton niche contains a single firm by construction, so it exhibits no within-niche variation in entry and cannot contribute to identifying the congestion relationship. Excluding singletons biases the sample toward congested niches. Specifically, if singletons disproportionately represent novel ideas that escape congestion, the remaining sample overstates the degree of crowding. The endogeneity discussion in Section 7.4 addresses this concern.

**Table 2:** Summary Statistics

Variable	Mean	Median	SD	Min	Max
<i>Panel A: Startups (N = 37,818)</i>					
Exited within 6 years	0.020	0.000	0.142	0.000	1.000
Total funding (million USD)	25.273	0.000	389.366	0.000	61900.120
Founded year	2014.962	2015.000	1.978	2012.000	2018.000
<i>Panel B: Ideas (N = 2,045)</i>					
Exit rate	0.029	0.000	0.088	0.000	1.000
Number of exits	0.378	0.000	1.626	0.000	39.000
Media mentions	751.896	120.333	2275.654	0.000	41723.875
Total entrants (all cohorts)	18.493	8.000	27.331	2.000	232.000
<i>Panel C: Idea-Years (N = 9,559)</i>					
Exit rate	0.025	0.000	0.127	0.000	1.000
Media mentions	924.091	160.000	3068.526	0.000	89469.000
Annual entrants	3.956	2.000	4.921	1.000	53.000
Log(annual entrants)	0.908	0.693	0.899	0.000	3.970

Notes: Panel A reports statistics at the startup level. Exited within 6 years equals one if the firm completed an IPO or was acquired within six years of founding. Total funding is the cumulative amount raised across all recorded funding rounds, expressed in millions of USD. Panel B reports statistics at the product-market niche level, where each niche is a cluster of startups assigned to the same idea. Media mentions is the count of news articles referencing the niche in US national online news sources and serves as our proxy for perceived idea quality. Total entrants counts all startups founded between 2012 and 2018 assigned to the niche. Panel C reports statistics at the idea-year level, which is the unit of observation in the main regressions. Annual entrants is the number of startups founded in a given year within a niche. Exit rate is the fraction of annual entrants that exited within six years.

founding-year cohort enters in 2018 and our Crunchbase snapshot is from mid 2025. Furthermore, a shorter window would increase right-censoring of exits that occur later in a startup’s life since an exit within six years is already considered quite rapid. The resulting sample spans 37,818 startups across roughly 2,045 niches. Panel A of [Table 2](#) reports startup-level summary statistics. Appendix [Table D1](#) provides step by step sample attrition information.

## 7.2 Defining Ideas from AI Ad-Keyword Portfolios

The model’s ideas are product-market niches in which startups compete for the same customers with similar products. Measuring these niches requires a method that can assign every startup in the sample to a group defined by demand-side substitutability. DoorDash (restaurant meal delivery) and Instacart (grocery delivery) serve different buyer intents and belong in different niches despite both being delivery platforms. Our method must make this distinction reliably at scale.

Published methods for measuring product similarity rely on detailed 10-K product descriptions standardized by SEC disclosure rules.<sup>8</sup> These methods do not extend to startups. Startup descriptions are short, marketing-oriented, and lack regulatory standardization. One description in our data read “Application: Uber for dog walking. Hardware: Fitbit for Dogs,” referring to a dog-walking marketplace plus a pet activity tracker. Our implementations of various NLP methods matched this company with *ride-sharing* or *wearable devices* startups and routinely grouped Instacart with DoorDash.<sup>9</sup>

Rather than forcing similarity out of thin textual descriptions, we use large language models (LLMs) to infer the concrete product and use-case details that are implicit in short descriptions. In our work, we ask LLM agents embedded in the explicit role of a Google AdWords specialist in the year companies were founded to represent idea boundaries by the set of customer search intents the firm would rationally pay to reach. Thus, we simulate Google Ads-style keywords the firm would bid on if it were trying to acquire customers in its target niche. In this way, we characterize competition as a revealed preference object controlled by the firm.<sup>10</sup>

One may worry about a “garbage-in, garbage-out” scenario. The concern is that if company descriptions are too broad, the generated keywords could be just as noisy as direct text-similarity scores. Generative models, however, excel at filling in plausible latent details from sparse inputs, predicting the concrete product attributes and customer segments implied by an abbreviated description. Recent work validates this capability in economic contexts. For instance, LLM agents recover theoretically predicted bidding behavior in auction experiments (Manning *et al.*, 2024) and replicate qualitative patterns from classic behavioral experiments (Filippas *et al.*, 2024). Our setting is less demanding, since we ask the model to simulate a marketing exercise with a well-defined right answer rather than to exhibit equilibrium behavior.

Another concern is that LLM-generated keywords reflect the model’s training biases rather than actual startup positioning. The keywords need only preserve relative competitive proximity between firms, not recover true advertising bids. We validate this relative accuracy

---

<sup>8</sup>Hoberg & Phillips (2010) develop a text-based approach for calculating pairwise similarity scores between firms by analyzing product descriptions from company 10-K forms (see also Hoberg & Phillips, 2016). Pellegrino (2025) further maps this similarity matrix into a matrix of cross-price demand elasticities inside a scalable hedonic demand system to back out markups and market power for the universe of U.S. public firms.

<sup>9</sup>Tuning a clustering algorithm’s resolution parameter can eventually separate such companies. Doing so, however, raises precision at the expense of drastically lower recall; that is, separating companies that should be kept together.

<sup>10</sup>Using Google Ads data directly is infeasible for our sample size. First, most startups do not run paid marketing campaigns. Second, even after negotiating research-use, historical bidding data proved prohibitively expensive to acquire.

using a hand-labeled benchmark, which confirms that the ordinal ranking holds.

To generate keywords, we use ExpectedParrot (<https://www.expectedparrot.com>), an open-source platform for conducting large-scale surveys with LLM agents. Figure D1a shows the wording of the marketing keyword question we ask and Figure D1b shows the traits we gave to the LLM agents. In reality, all keywords do not have the same level of importance in capturing a customer intent (e.g., for DoorDash, ‘food delivery’ has a higher reach than ‘restaurant takeout’). Therefore, we ask a second question from the same LLM agent about what fraction of its marketing budget would it allocate to each of the ten keywords (see Figure D3 for details). For instance, the keywords and corresponding fraction of budget allocated to each for DoorDash are ‘food delivery’ (0.15), ‘restaurant delivery’ (0.15), ‘takeout delivery’ (0.08), ‘meal delivery’ (0.12), ‘delivery near me’ (0.13), ‘late night delivery’ (0.07), ‘online food order’ (0.08), ‘local food delivery’ (0.10), ‘delivery service’ (0.07), and ‘restaurant takeout’ (0.05).

**Clustering:** We use these weighted keyword portfolios to measure pairwise similarity between firms via a weighted Jaccard index:

$$J(i, j) = \frac{\sum_k \min\{w_{ik}, w_{jk}\}}{\sum_k \max\{w_{ik}, w_{jk}\}},$$

where  $w_{ik}$  is the normalized budget weight on keyword  $k$  in firm  $i$ ’s portfolio. The weighted Jaccard respects the ordinal importance of keywords. DoorDash and Instacart, for example, share zero keywords after normalization ( $J = 0$ ) despite both being delivery platforms, because their portfolios target distinct buyer intents (restaurant meals vs. groceries). DoorDash and Zoomer, a smaller restaurant delivery startup, share five keywords with closely aligned weights ( $J \approx 0.40$ ) and land in the same cluster. The natural alternative, cosine similarity on keyword vectors, assigns positive similarity to any pair of firms and would blur niche boundaries. The Jaccard index is zero whenever two firms share no keywords, so the DoorDash-Instacart separation at  $J = 0$  is exact rather than a matter of choosing a similarity threshold.

We convert these pairwise similarities into a mutual  $k$ -nearest-neighbor graph, meaning each firm nominates its  $k$  most similar peers, and we retain a link only when both firms nominate each other. We then partition the graph using the Leiden algorithm with a Constant Potts Model (CPM) objective, where the resolution parameter  $\gamma$  controls how fine-grained the partition is (a higher  $\gamma$  produces more, smaller niches).<sup>11</sup> The two tuning

---

<sup>11</sup>Leiden improves on the widely used Louvain algorithm by guaranteeing that every detected community is internally connected, eliminating the poorly connected clusters Louvain can produce. The CPM objective does not scale the resolution with graph size, so cluster granularity remains interpretable as the sample

parameters— $k$  and  $\gamma$ —govern graph density and cluster granularity, respectively. We select  $k = 180$  and  $\gamma = 0.028$  by maximizing agreement with a hand-labeled benchmark of 268 startups spanning 65 product-market labels.<sup>12</sup> On this benchmark, B<sup>3</sup> precision (the share of predicted cluster-mates that truly belong to the same niche) is 0.91, and B<sup>3</sup> recall (the share of true niche-mates that the algorithm places together) is 0.81. Adjusted mutual information, which measures how well the algorithm’s groupings match the hand labels after correcting for chance agreement, is 0.86. For further implementation details, see [Appendix C.1](#).

This mapping defines the unit of competition for our empirical analysis, where entry is the inflow of startups assigned to an idea in a founding year. After dropping singleton clusters and niches with no IPO or acquisition, the final sample contains 37,818 startups across 2,045 niches. As described in [Table 2](#), the median niche has 8 entrants, the largest has 232, and the size distribution is right-skewed, consistent with the model’s prediction that a few high-quality ideas attract many entrants. Exits are rare. Only 2% of startups complete an IPO or acquisition within six years of founding, and the median niche has an exit rate of zero. This rarity is consistent with the high failure rates documented in the entrepreneurship literature and provides substantial cross-niche variation in the outcome we use to test the model’s congestion predictions.

### 7.3 Measuring Idea Quality: Media Attention as a Public Signal

The model’s comparative statics rank ideas by quality  $q$  and predict that higher-quality ideas attract more entrants. Testing these predictions requires an observable correlate of  $q$  that varies across niches. To proxy idea quality, we track media attention via online news coverage. Economically, attention acts as a public signal of salience and perceived promise, helping to coordinate the beliefs of entrepreneurs, investors, and workers. While media attention is not realized performance, it offers a valuable external signal of how an idea is perceived outside the platform providing the startup descriptions.

We derive this measure from Media Cloud, a platform API that analyzes news coverage from curated collections ([Roberts \*et al.\*, 2021](#)). Because Media Cloud documents the underlying sources in each collection, we can hold the set of news producers fixed over time. Given a search query, the API returns time-stamped matches we aggregate into story counts.

Connecting this proxy to our niche definitions involves building idea-level search queries from the same demand-side language used in clustering. For each niche, we pool the keyword portfolios of all member firms, sum the budget weights across firms that share a normalized

---

grows.

<sup>12</sup>This is an unsupervised procedure in the sense that no labels enter the clustering objective. The labeled datapoints are used only to select tuning parameters.

phrase, and retain the 24 highest-weight phrases.<sup>13</sup> Aggregated to the year level, the resulting “story counts” produce an idea-by-year measure of media attention merged back to our entry panel. The median niche receives 120 mentions, the distribution is heavily right-skewed, with a mean of 752 (Panel B of [Table 2](#)).

The model’s predictions are stated in terms of true idea quality  $q$ , but the regressions use media mentions  $m$  as a proxy. If  $m$  is a monotone increasing function of  $q$ , the sign of the entry-quality relationship is preserved. Whether the log-log specification recovers the theoretical unit elasticity depends on the log-elasticity of media mentions with respect to true quality,  $\varepsilon_h = d \log m / d \log q$ . When  $\varepsilon_h < 1$  (the press under-covers the best ideas), the measured entry-quality gradient is steepened; when  $\varepsilon_h > 1$  (the press amplifies top ideas), it is attenuated. [Appendix C.2](#) derives these conditions formally and discusses implications for interpreting the regression coefficients.

## 7.4 Empirical Results

The model generates three testable predictions. First, entry increases in idea quality. Second, per-firm success declines with entry, and the elasticity of success with respect to entry lies between  $-1$  and  $0$  ([Equation 6](#)). Third, idea-level success is concave in entry, implying diminishing returns that the decentralized equilibrium fails to internalize ([Proposition 4](#)). We examine each prediction in turn, first descriptively and then with regression evidence. All specifications are cross-sectional or absorb year fixed effects and we interpret the results as informative correlations, not causal estimates.


We start with a visual description of the key results in [Figure 2](#). The top figure shows that higher-quality ideas attract more entry. Moving from the lowest to the highest quality decile, the mean number of entrants nearly triples. Because the figure weights each idea-year by its number of entrants, the y-axis reports the competition experienced by a randomly drawn firm rather than the average number of entrants per idea.

The middle figure turns to per-startup success probabilities. The relationship with idea quality is noisy and statistically indistinguishable from flat across most of the distribution. Even in the top quality decile, the six-year exit probability remains on the order of a few percent. The absence of a clear quality gradient is consistent with the model’s prediction that congestion compresses ex ante differences in idea quality into similar ex post success rates. It is possible, however, that media mentions may simply be too noisy, or that the sparse

---

<sup>13</sup>The weight aggregation acts as a revealed-preference filter so that phrases that many firms in the niche independently prioritize rise to the top. These phrases are quoted and joined into OR queries submitted to Media Cloud. Researchers interested in using Media Cloud for large-scale API calls are encouraged to coordinate with the team to avoid causing a service interruption.

**Figure 2:** Market Entry and Success by the Perceived Quality of Product-Market Ideas



Notes: This figure plots the relationship between product-market idea quality and both competition and success at the idea-cohort level, where the unit of observation is a product-market idea-firm entry year. Idea quality deciles are constructed from the within-year standardized media coverage of the product-market in the calendar year prior to founding, using entrant weights so that each startup contributes equally. The top figure reports the firm-weighted mean number of entrants by quality decile. The middle figure reports the firm-weighted mean probability that a startup in that decile exits within six years. The bottom figure reports the probability that an idea-year produces at least one successful exit within six years; each idea-year is weighted equally because the quantity of interest is a property of the idea rather than of the firm. Shaded regions report 95 percent confidence intervals.

**Table 3:** Entry, Quality of Ideas, and Market Success

DV:	Log(Entry)		Fraction Exited within 6 Years	
	3-1	3-2	3-3	3-4
Log(Media Mentions)	0.098*** (0.006)	0.088*** (0.017)		0.000 (0.001)
Log(Media Mentions) <sup>2</sup>		0.001 (0.002)		
Log(Annual Entry)			-0.008*** (0.002)	-0.008*** (0.002)
<i>N</i>	9,559	9,559	9,559	9,559
Mean of DV	1.321	1.321	0.025	0.025
Year FE	X	X	X	X

Notes: This table reports linear regressions at the idea-year level, where the unit of observation is an idea-firm entry year. The dependent variable in Columns 3-1 and 3-2 is the log number of startups entering that product-market in a given year. The dependent variable in Columns 3-3 and 3-4 is the fraction of startups in that idea-year that achieve an exit within six years of founding. Idea quality is measured by the within-year standardized media coverage of the product-market in the calendar year prior to founding. Column 3-2 adds a quadratic term in log media mentions to test for departures from log-linearity in the entry-quality relationship. Standard errors clustered by ideas are reported in parentheses. Statistical significance is \*(10%), \*\*(5%), or \*\*\*(1%).

binary nature of the outcome may leave insufficient statistical power to detect a monotone relationship even if one exists. In addition, the slight uptick in the top decile may reflect the tail of the quality distribution outrunning congestion.

The bottom figure reports the probability that an idea produces at least one successful exit within six years. The probability rises monotonically with quality, roughly tripling from the bottom to the top decile. Per-firm success is flat across these deciles, yet higher-quality ideas are more likely to succeed because they attract more entrants, each with an independent chance of exiting successfully. The pattern suggests that the media-based proxy carries signal at the idea level even when congestion compresses differences in per-firm outcomes.

Table 3 provides regression evidence that mirrors the descriptive patterns in Figure 2. In Column 3-1, we find that higher-quality ideas attract more startups. The coefficient on log(Media Mentions) implies that a 10% increase in media mentions is associated with roughly a 1% increase in the number of entrants targeting that idea within a founding-year cohort. The strength of this relationship is modest in magnitude, but consistent with the notion that ex-ante more promising ideas draw in more entrepreneurs.

Column 3-2 tests the functional form of the entry-quality relationship by adding a quadratic term, log(Media Mentions)<sup>2</sup>. The model predicts that entry is proportional to idea quality ( $n(q) = q/q_0$ ), which implies a linear log-log relationship between entry and

quality. A significant quadratic term would indicate departure from this prediction. The estimated quadratic term is small and statistically indistinguishable from zero, consistent with the proportional entry prediction.

Columns 3-3 and 3-4 turn to per-firm success. Higher entry is associated with lower exit probabilities. In Column 3-4, a 10% increase in entry is associated with a 0.08 percentage point decrease in the exit rate, a 3.2% decline relative to the mean. Idea quality itself has no discernible residual association with success once entry is controlled for. The model’s equalization logic predicts this pattern. Free entry into higher-quality ideas compresses differences in per-firm success across the idea distribution. Columns 3-1 through 3-4 together show both halves of the mechanism: 1) quality attracts entry, and conditional on entry, 2) quality has little residual association with exit probabilities.

Table 4 estimates the elasticity of success with respect to entry. We collapse to the product-market idea level, so that each observation is a niche aggregated across all sample years. The cross-niche regression identifies the model’s system-wide  $\eta_E^{q_0}$  under the assumption that niches behave as separate free-entry equilibria, sharing the model’s structural parameters but differing in scale ( $E_i$ , the mass of entrepreneurs entering each niche), with flows between niches limited by skill specialization.<sup>14</sup> Under this assumption, cross-sectional variation in  $\log n_i$  traces out the comparative static of  $\log q_{0i}$  with respect to  $\log E_i$  that the model derives within a single system, and the elasticity  $\eta$  governs whether marginal entrants pursue novel ideas ( $\eta \rightarrow 0$ ) or crowd into ideas already targeted by incumbents ( $\eta \rightarrow -1$ ).

Table 4 reports three estimators of the congestion elasticity. The first two estimates use the full sample of 2,045 niches, and the third one restricts to the 458 niches with at least one exit within six years of founding, since the log of the exit rate is undefined when no exits occur. Column 4-1 reports baseline OLS estimates in levels. The coefficient of  $-0.012$  implies that a 10% increase in total entry is associated with a 0.1 percentage point reduction in the exit rate, a 3.9% decline relative to the sample mean of 2.9%. Column 4-2 reports our preferred Poisson estimator, which models the count of exits with total entry as exposure and avoids the log restriction. The estimated elasticity is  $-0.18$ , meaning a 10% increase in the mass of entrepreneurs entering technology entrepreneurship is associated with a 1.8% decrease in the per-firm probability of success. Column 4-3 shows the log-log OLS estimates on the restricted subsample, yielding an elasticity of  $-0.86$ . We interpret the larger magnitude as driven by two distinct forces. First, conditioning on at least one within-six-year

---

<sup>14</sup>Were skills perfectly shared across niches, free entry would equalize per-firm success at a single  $q_0$  across all niches, and the cross-niche regression slope would be zero regardless of the underlying crowding parameter. Partial skill-sharing biases the cross-niche slope toward zero, so the full-sample Poisson estimate is a lower bound on the true  $|\eta_E^{q_0}|$ . The slope being nonzero in the data is therefore evidence that segmentation is at least partially binding.

**Table 4:** Elasticity of Success Probability to Entry

Estimator:	OLS Levels	Poisson	OLS Log-Log
	4-1	4-2	4-3
Log(Total Entry)	-0.012*** (0.002)	-0.182** (0.089)	-0.861*** (0.024)
Log(Media Mentions)	0.003** (0.001)	0.102*** (0.038)	0.028** (0.013)
Sample	Full Sample	Full Sample	Successful Niches
$N$	2,045	2,045	458
Mean of DV	0.029	0.378	-2.644

Notes: This table reports cross-sectional regressions at the product-market idea level, where each observation is a product-market cluster aggregated across all sample years. Column 4-1 reports OLS in levels, with dependent variable equal to the number of exits divided by total entrants. Column 4-2 reports a Poisson regression of the count of exits with total entry as the exposure variable. In this model, the coefficient on Log(Total Entry) is the elasticity directly. Column 4-3 reports OLS in log-log, with dependent variable equal to log exit rate, restricted to the 458 niches with at least one observed exit within six years since. This specification drops niches with no exits within six years. Robust standard errors are reported in parentheses. Statistical significance is \*(10%), \*\*(5%), or \*\*\*(1%).

exit selects niches whose realized exit count is positive, which is mechanically increasing in the number of entrants. This selection channel inflates the magnitude of the estimate, while the segmentation bias from partial skill-sharing pushes in the opposite direction; we take the selection channel to dominate on this restricted subsample. Second, the model’s success-probability concept is most directly identified on the most viable ideas. We treat  $-0.18$  as the conservative full-sample estimate and  $-0.86$  as an upper bound on the congestion elasticity for viable ideas.

The two estimates jointly place the true congestion elasticity in the range  $[-0.86, -0.18]$ , both within the theoretical  $[-1, 0]$  continuum the model identifies. The upper bound on viable ideas implies that the large majority of marginal entry crowds into ideas already targeted by incumbents, generating business-stealing externalities rather than expanding the frontier of explored opportunities. Even the lower bound on the full sample places a non-trivial share of marginal entry on the intensive margin.

This estimate is correlational, and three sources of endogeneity deserve explicit discussion. First, media mentions and entry are jointly determined. In particular, a technological breakthrough in year  $t-2$  can generate press coverage in  $t-1$  and attract entrepreneurs in  $t$ . Our quality proxy uses media mentions from the calendar year prior to founding, which mitigates the most direct form of reverse causality (entry causing press) but does not eliminate

confounding from common shocks. Second, unobserved niche characteristics that raise entry and lower per-firm success (low barriers, commoditized technology attracting many undercapitalized entrants) would inflate the estimated  $|\eta|$ . Third, as derived in [Appendix C.2](#), the log-elasticity of media mentions with respect to true quality determines whether the proxy introduces its own bias in the entry regression, though this channel operates on the quality coefficient rather than directly on  $\eta$ .

An ideal instrument would provide exogenous variation in the number of entrants per idea that is uncorrelated with idea quality or niche-level unobservables, such as a supply-side shock that steers engineers toward entrepreneurship in some niches but not others. The AWS design below moves in that direction by using an externally timed cost and tooling shock that should matter more in niches already oriented toward cloud-heavy production.

Tables [3](#) and [4](#) establish the per-firm congestion mechanism. The bottom figure of [Figure 2](#) showed descriptively that higher-quality ideas are more likely to produce at least one winner. The model’s efficiency prediction in [Proposition 4](#) turns on the shape of the idea-level success function. If individual startup success probabilities are i.i.d. with probability  $p$ , the probability that at least one of  $n$  entrants succeeds is  $P(n) = 1 - (1 - p)^n$ . This function is concave in  $n$ , thus the decentralized allocation  $n(q) = q/q_0$  is linear in idea quality, while the socially optimal allocation  $n^*(q)$  is concave. The divergence between the two is steepest for the highest-quality ideas, which attract the most entry but yield the smallest marginal contribution to idea-level success per additional founder.

[Table 5](#) tests this prediction at the idea level. The dependent variable in [Columns 5-1](#) and [5-2](#) is an indicator equal to one if any startup pursuing that idea achieved a successful exit within six years. [Column 5-1](#) estimates that each additional entrant is associated with a half percentage point increase in the probability that the idea produces at least one winner, a meaningful effect against a baseline success rate of 22 percent. [Column 5-2](#) adds a quadratic term to test for the concavity. Since  $P(n)$  is concave in  $n$  rather than in  $\log n$ , a quadratic in levels directly maps to the curvature of the theoretical success function.<sup>15</sup> The quadratic coefficient on total entry squared is negative and precisely estimated, confirming diminishing returns to entry at the idea level. [Appendix Figure D4](#) provides a visual test showing the binned scatterplot of idea-level success rates against entry tracks the theoretical prediction  $P(n) = 1 - (1 - \hat{p})^n$  closely for small to moderate  $n$ , with the most crowded niches falling

---

<sup>15</sup>We specify the quadratic in levels rather than logs because the theory’s concavity prediction applies to  $P(n) = 1 - (1 - p)^n$  as a function of  $n$ , not of  $\log n$ . The second derivative  $P''(n) = -[\ln(1 - p)]^2(1 - p)^n$  is negative for all  $n$ , confirming strict concavity. However, expressing the same function as  $Q(x) = P(e^x)$  where  $x = \log n$  yields  $Q''(x) > 0$  whenever  $n < 1/p$ . For  $p \approx 0.02$ , this threshold is 50, which exceeds the median niche size in our data. A quadratic in logs would therefore produce a positive second-order coefficient that reflects this coordinate transformation, not a departure from diminishing returns.

**Table 5:** Idea-Level Success and Diminishing Returns to Entry

DV:	Any Exit within 6 Years		Number of Exits
	5-1	5-2	5-3
Total Entry	0.005*** (0.000)	0.007*** (0.001)	0.023*** (0.006)
Log(Media Mentions)	0.012*** (0.004)	0.009** (0.004)	0.006 (0.010)
Total Entry <sup>2</sup>		-0.000*** (0.000)	
<i>N</i>	2,045	2,045	2,045
Mean of DV	0.224	0.224	0.378

Notes: This table reports cross-sectional regressions at the product-market idea level, where each observation is a product-market cluster aggregated across all sample years. The dependent variable in Columns 5-1 and 5-2 is a binary indicator equal to one if any startup pursuing that idea achieved an IPO or acquisition within six years of founding. The dependent variable in Column 5-3 is the count of startups on that idea that achieved such an exit. Total Entry is the number of startups entering the product-market across all sample years, measured in levels. Column 5-2 adds a quadratic in total entry to test for concavity of the idea-level success function. Log(Media Mentions) controls for idea quality. Robust standard errors are in parentheses. Statistical significance is \*(10%), \*\*(5%), or \*\*\*(1%).

slightly below the curve.

Column 5-3 shifts to the intensive margin. Under the strict winner-take-all assumption, the number of winners should be approximately invariant to entry. The estimated coefficient, however, suggests that each additional entrant adds roughly 0.02 exits per idea. At the mean niche size of 18 entrants, this implies about 0.36 exits, close to the sample mean of 0.38. Doubling entry from the mean would add another 0.36 exits, nearly doubling the count but still implying strong diminishing returns given the concavity established in Column 5-2. The strict single-winner structure is rejected, but the data are consistent with a market that accommodates a small number of winners even as entry grows substantially.

Taken together, the three tables paint a consistent picture. Entry responds to idea quality, per-firm success declines with entry, and idea-level success is concave in the number of entrants. The estimated elasticity of  $-0.18$  and the negative quadratic in Column 5-2 jointly suggest that the highest-quality ideas attract the most entry while yielding the smallest marginal contribution to idea-level success per additional founder. This is the “best-idea trap” predicted by the model, and it is the pattern that drives the welfare wedge in Proposition 4.

## 7.5 Externally Timed AWS Cost Shocks

The cross-sectional estimates above establish the congestion mechanism, but they remain vulnerable to the concern that entry and idea quality are jointly determined. To introduce plausibly exogenous variation in entry, we study a set of AWS cost and tooling shocks that differentially lowered the cost of building cloud-intensive startups. Amazon Web Services (AWS) is the dominant cloud-computing platform used by startups to rent computing power, storage, databases, and backend infrastructure rather than building those capabilities in-house. For young software firms, improvements in AWS can matter directly for fixed and quasi-fixed startup costs: cheaper infrastructure lowers the cost of experimentation, while new managed tools reduce the engineering effort needed to launch a product. For each product-market niche, we measure pre-shock exposure using the share of pre-period startup descriptions that load on cloud, serverless, API, database, hosting, and related terms. We then estimate difference-in-differences specifications of the form  $\text{post}_t \times \text{exposure}_j$  with niche and year fixed effects.

The closest antecedent for this identification strategy is [Ewens \*et al.\* \(2018\)](#), who use Amazon’s 2006 introduction of cloud computing as a discrete reduction in the cost of starting cloud-exposed ventures and document a shift toward “spray and pray” venture-capital investing in treated industries.<sup>16</sup> We complement their investor-side analysis with an entrepreneur-side analysis at a finer niche-level definition.

The broadest design combines the March 2014 AWS price cuts with the November 2014 launches of Lambda and Aurora, with the post period beginning in 2015. That bundled “cloud wave” captures a general decline in the cost of standing up cloud-native startups. The narrower shocks isolate more specific technical margins. Lambda made “serverless” computing easier by allowing firms to run code without managing their own servers. API Gateway simplified the task of exposing and managing application-programming interfaces, which is particularly relevant for API-first and backend software businesses. [Table 6](#) reports the entry and exit-rate estimates for the broad cloud-wave design and for the narrower Lambda and API Gateway shocks. The broad cloud-wave shock yields the clearest result: more exposed niches experience significantly more entry after the shock, with a coefficient of 0.0920, implying roughly a 9 percent increase in entrants for a fully exposed niche relative to an unexposed one. By contrast, the corresponding exit-rate coefficient is small and statistically indistinguishable from zero. The same basic pattern appears for Lambda and

---

<sup>16</sup>They define treatment at the level of 26 VentureSource industry-segments and study investor-side portfolio outcomes including initial round size, board seats, founder experience, and failure rates. Our Crunchbase sample begins in 2012, so we cannot study their 2006 event and instead exploit the 2014–16 wave of AWS cost reductions and managed-service launches. AWS had been the dominant cloud platform for nearly a decade by 2014, so the 2014 shock is smaller than the 2006 introduction.

**Table 6:** AWS Cost and Tooling Shocks: Entry and Exit-Rate Effects

	Log entrants	Exit rate by 6 years
<i>Panel A: 2014 Cloud Wave (post = 2015)</i>		
Post $\times$ exposure	0.0920** (0.0415)	-0.0075 (0.0092)
Observations	9,132	9,132
Pre-trend test ( $p$ -value)	0.313	
<i>Panel B: AWS Lambda (post = 2015)</i>		
Post $\times$ exposure	0.1271** (0.0595)	0.0017 (0.0112)
Observations	9,132	9,132
Pre-trend test ( $p$ -value)	0.492	
<i>Panel C: API Gateway (post = 2016)</i>		
Post $\times$ exposure	0.1409*** (0.0462)	-0.0027 (0.0116)
Observations	9,401	9,401
Pre-trend test ( $p$ -value)	0.089	

Notes: This table reports niche-year regressions for three AWS-related cost and tooling shocks. Each panel corresponds to one shock definition, and the treatment is the interaction of a post-shock indicator with a niche-level pre-period exposure score constructed from startup text. The entry outcome is log entrants and the exit outcome is the niche-year exit rate within six years. All specifications include niche fixed effects, year fixed effects, and controls for log media mentions. Standard errors clustered by niche are reported in parentheses. The pre-trend row reports the  $p$ -value from a clustered joint Wald test that all pre-period event-study coefficients equal zero. Statistical significance is \*(10%), \*\*(5%), or \*\*\*(1%).

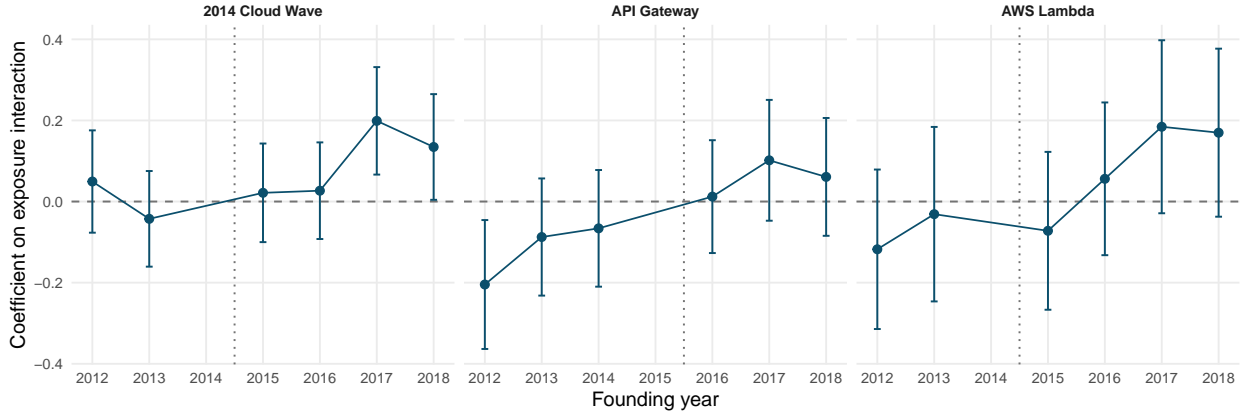
API Gateway.

Figure 3 reports the event-study evidence for the three AWS shock definitions in the main analysis. The cloud-wave and Lambda designs pass formal joint pre-trend tests comfortably, while API Gateway is only marginal on this dimension ( $p = 0.089$ ). A fourth shock based on Amazon ECS, which lowered the cost of deploying containerized applications, yields a comparable entry response but formally rejects flat pre-trends (see Table D2 and Figure D5).

Figure 4 reports the cloud-wave coefficient as we restrict the niche-level sample to those with at least  $N$  pre-period firms, for  $N \in \{1, 2, 3, 5\}$ . The coefficient strengthens monotonically with  $N$ , consistent with classical measurement-error attenuation in the exposure score for thinly-observed niches. We interpret this pattern as our main estimates in Table 6 showing a conservative lower bound.


Taken together, the AWS evidence is the reduced-form analogue of the paper’s main cross-sectional result that cost-reducing shocks expand startup formation more reliably than they improve realized niche success.

**Figure 3:** Event-Study Evidence for the Three Main AWS Shock Definitions



Notes: This figure plots event-study coefficients for the three AWS shock definitions in the main analysis, normalizing the year before the post period to zero within each panel. Treatment intensity is the niche-level pre-period exposure score for the corresponding shock. The dashed horizontal line marks zero and the dotted vertical line marks the start of the panel-specific post period.

**Figure 4:** Cloud-Wave Coefficient Across Pre-Period Support Thresholds



Notes: This figure plots the cloud-wave entry coefficient when the niche-level sample is restricted to niches with at least the minimum number of pre-period firms shown on the horizontal axis. The leftmost point ( $N = 1$ ) reproduces the headline estimate from [Table 6](#). Vertical bars are 95% confidence intervals based on standard errors clustered by niche.

## 7.6 A Second Setting: Thematic Crowding in Y Combinator

The Crunchbase analysis estimates congestion in a broad sample of US technology startups. A complementary test comes from a very different setting: Y Combinator’s public company directory. YC is not a population sample of startup entry. It is a selected accelerator portfolio, and batch composition reflects YC’s admissions choices as well as founder entry decisions. This difference is useful. The Crunchbase sample speaks to congestion across the entrepreneurial ecosystem, while the YC sample asks whether the same mechanism appears inside a curated institution where the aggregate scale of entry is partly chosen by the accelerator.


We collect 5,754 YC-backed companies from the public YC company directory, covering batches S05 through S26. For each company, we observe batch, industry, current YC status, and short business descriptions. We embed the concatenated one-line and long descriptions using OpenAI’s `text-embedding-3-small` model and cluster the resulting vectors into 50 thematic groups. The main YC crowding measure is the number of companies in the same embedding cluster and the same YC batch. The outcome is an indicator for whether the company is currently listed as inactive. Because recent batches are heavily right-censored, the crowding regressions restrict to batches from 2022 or earlier.

**Figure 5** shows two facts that connect the YC setting to the model. First, YC cohorts exhibit pronounced idea waves. Within-batch semantic similarity is relatively stable through the pre-2020 period and then rises sharply in recent cohorts, consistent with founders clustering around common technological opportunities, especially AI-related themes. Second, the relevant crowding margin is not the size of the YC batch as a whole but the number of close thematic peers funded at the same time. The residualized scatter in Panel (b) shows a positive relationship between cluster-batch crowding and inactivity after removing year effects.

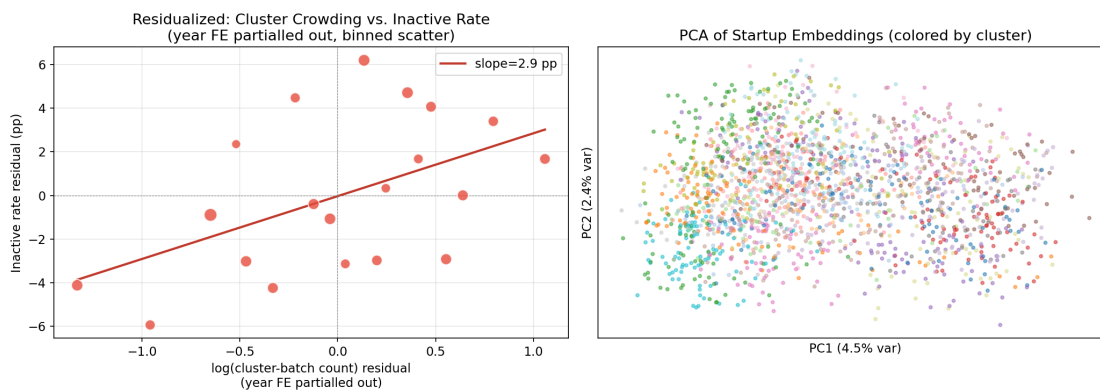
**Table 7** sharpens this distinction. Aggregate batch size is not associated with inactivity once cohort timing is controlled for. The coefficient on log batch size is close to zero with year fixed effects, and remains close to zero after adding industry controls. By contrast, the same-batch thematic count is positively associated with inactivity. A doubling of the number of close thematic peers in the same batch is associated with roughly a three percentage point increase in the probability of inactive status, and the estimate is stable after adding year, cluster, and industry fixed effects. The continuous peer-similarity measure is not statistically distinguishable from zero, suggesting that the economic margin is closer to the count of firms occupying the same niche than to semantic proximity alone.

The two empirical settings therefore play different roles. The Crunchbase analysis estimates the paper’s central congestion elasticity in a broad sample with demand-side niche

**Figure 5:** Thematic Convergence and Same-Batch Crowding in Y Combinator



(a) Within-batch semantic similarity



(b) Residualized cluster-batch crowding and embedding structure

Notes: Panel (a) reports the average pairwise cosine similarity among company description embeddings within each YC batch. Panel (b) plots the relationship between inactivity and log cluster-batch count after residualizing both variables on year fixed effects, and shows a PCA projection of company embeddings colored by cluster. The sample is the public YC company directory collected in March 2026.

**Table 7:** Batch Scale, Thematic Crowding, and YC Inactivity

	Batch scale		Thematic crowding		
	(1)	(2)	(3)	(4)	(5)
Log(batch size)	-0.003 (0.014)	0.001 (0.014)			
Log(cluster-batch count)			0.037*** (0.009)	0.028*** (0.009)	
Peer cosine similarity					0.000 (0.011)
Year fixed effects	Yes	Yes	Yes	Yes	Yes
Industry fixed effects	No	Yes	No	Yes	No
Cluster fixed effects	No	No	No	No	No
Sample	All batches	All batches	Mature batches	Mature batches	Mature batches
N	5,753	5,753	3,793	3,793	3,793

Notes: This table reports company-level linear probability models using the public YC company directory. The dependent variable is an indicator for current inactive status. Columns (1) and (2) test whether aggregate YC batch size predicts inactivity. Columns (3) and (4) replace aggregate batch size with log cluster-batch count, defined as the number of companies in the same embedding cluster and YC batch. Column (5) uses average cosine similarity to same-cluster same-batch peers as a continuous alternative crowding measure. Standard errors are clustered by batch. The mature-batch sample restricts to companies in batches from 2022 or earlier. Statistical significance is \*(10%), \*\*(5%), or \*\*\*(1%).

definitions and six-year exit outcomes. The YC analysis shows the same mechanism in an institutional microcosm with a different sample, text representation, and outcome. Together they point to the same object: congestion is generated by dense entry into a specific product-market opportunity, not by entrepreneurial activity in the aggregate.

## 8 Conclusion

This paper develops a model of entry into technology entrepreneurship, treating an entrepreneurial cluster as three interconnected markets for venture capital, skilled labor, and product markets. The model generates four results that flow from a single mechanism. Because entrepreneurs observe the same opportunities and enter freely, equilibrium success probabilities equalize across all pursued ideas, generating a congestion externality that individual entrepreneurs do not internalize. This equalization implies that the decentralized equilibrium features excessive entry relative to the social optimum, that the excess entry concentrates on the highest-quality ideas when startups face idiosyncratic execution risk, and that engineer wages serve as a sufficient statistic for total system output. The severity of the distortion is governed by a single parameter, the elasticity of startup success probability with respect to the number of entrants. The model also shows that the inefficiency can be lessened by public R&D that expands the set of usable ideas.

We provide evidence that is broadly consistent with these predictions. The evidence comes from two complementary settings. In a broad Crunchbase sample, product-market entry predicts lower per-firm success and concave idea-level success, and externally timed AWS cost shocks increase entry into exposed niches without raising downstream exit outcomes. In the YC sample, aggregate batch scale does not predict inactivity, but same-batch thematic crowding does. Most of the empirical analysis still relies on proxies for idea quality and on cross-sectional variation that is informative about mechanisms rather than fully causal. The AWS designs add externally timed entry shifters, but they remain reduced-form and do not isolate a single structural elasticity of congestion. Establishing whether congestion causally reduces startup success probabilities with a sharper design would require exogenous variation in entry that is both stronger and more tightly mapped to specific product markets.

Our model is static and treats idea quality as common knowledge. These choices isolate congestion and keep the welfare analysis tractable, but they also point to natural extensions. Incorporating learning and staged experimentation would allow entrepreneurs to update beliefs about idea quality after entry, potentially amplifying or mitigating the congestion externality depending on how signals are correlated across entrants. Endogenizing the supply of ideas would connect congestion to the incentive to create new opportunities, since entrepreneurs who anticipate crowding on existing ideas face a higher return to generating genuinely novel ones. Similarly, introducing intellectual property protection would alter the mapping between idea quality and entry, since patents or trade secrets can limit the number of entrants per idea and change the equilibrium allocation.

The model's most distinctive policy implication is a market-failure justification for public R&D that differs from standard knowledge-spillover arguments. Expanding the supply of economically distinct ideas, through basic research, technology transfer, or targeted funding, directly reduces congestion by shifting marginal entry from the intensive margin to the extensive margin. At the same time, wages provide a tractable sufficient statistic for monitoring total system output.

For entrepreneurs and investors, the key implication is that entry density is a first-order determinant of risk. The most popular ideas in a cohort may systematically underperform because they attract the most competition, while less visible ideas near the viability threshold offer superior risk-adjusted returns. The ride-sharing market of the early 2010s, where dozens of well-funded startups pursued the same idea and most failed, is one illustration. The mechanism the model identifies is general. Wherever entrepreneurs can observe the same opportunities, congestion will compress differences in idea quality into similar outcomes.

## References

- ACEMOGLU, DARON, AKCIGIT, UFUK, BLOOM, NICHOLAS, & KERR, WILLIAM R. 2013. Innovation, reallocation and growth. *Working Paper*.
- AGHION, PHILIPPE, & HOWITT, PETER. 1992. A model of growth through creative destruction. *Econometrica*, **60**(2), 323.
- AGRAWAL, AJAY, GANS, JOSHUA S., & GOLDFARB, AVI. 2019. Exploring the impact of artificial Intelligence: Prediction versus judgment. *Information economics and policy*, **47**, 1–6.
- AGRAWAL, AJAY, GANS, JOSHUA S., & STERN, SCOTT. 2021. Enabling Entrepreneurial Choice. *Management Science*, **67**(9), 5510–5524.
- AGRAWAL, AJAY, CAMUFFO, ARNALDO, GAMBARDELLA, ALFONSO, GANS, JOSHUA S., SCOTT, ERIN L., & STERN, SCOTT (eds). 2026. *Bayesian Entrepreneurship*. Cambridge, MA: The MIT Press.
- AKCIGIT, UFUK, & KERR, WILLIAM R. 2010. Growth through heterogeneous innovations. *Working Paper*.
- BLANCHFLOWER, DAVID G, & OSWALD, ANDREW J. 1998. What makes an entrepreneur? *Journal of Labor Economics*, **16**(1).
- BRYAN, KEVIN A., HOFFMAN, MITCHELL, & SARIRI, AMIR. 2026. Information Frictions and Employee Sorting between Start-ups. *American economic journal. Applied economics*, **18**(2), 369–406.
- ELOUNDOU, TYNA, MANNING, SAM, MISHKIN, PAMELA, & ROCK, DANIEL. 2024. GPTs are GPTs: Labor market impact potential of LLMs. *Science (American Association for the Advancement of Science)*, **384**(6702), 1306–1308.
- EVANS, DAVID S, & JOVANOVIC, BOYAN. 1989. An estimated model of entrepreneurial choice under liquidity constraints. *Journal of Political Economy*, **97**(4), 808–827.
- EWENS, MICHAEL, NANDA, RAMANA, & RHODES-KROPF, MATTHEW. 2018. Cost of experimentation and the evolution of venture capital. *Journal of financial economics*, **128**(3), 422–442.
- FILIPPAS, APOSTOLOS, HORTON, JOHN J., & MANNING, BENJAMIN S. 2024. Large Language Models as Simulated Economic Agents: What Can We Learn from Homo Sili-cus? *Pages 614–615 of: Proceedings of the 25th ACM Conference on Economics and Computation*. New York, NY, USA: ACM.
- GANS, JOSHUA S., STERN, SCOTT, & WU, JANE. 2019. Foundations of entrepreneurial strategy. *Strategic management journal*, **40**(5), 736–756.
- GROMB, DENIS, & SCHARFSTEIN, DAVID. 2002. Entrepreneurship in equilibrium. *Working Papers*.
- HERRANZ, NEUS, KRASA, STEFAN, & VILLAMIL, ANNE P. 2013. Entrepreneurs, risk aversion and dynamic Firms. *Working Paper*.
- HOBERG, GERARD, & PHILLIPS, GORDON. 2010. Product Market Synergies and Competition in Mergers and Acquisitions: A Text-Based Analysis. *The Review of financial studies*,

- 23**(10), 3773–3811.
- HOBERG, GERARD, & PHILLIPS, GORDON. 2016. Text-Based Network Industries and Endogenous Product Differentiation. *The Journal of political economy*, **124**(5), 1423–1465.
- HOLMES, THOMAS J, & SCHMITZ JR, JAMES A. 1990. A theory of entrepreneurship and its application to the study of business transfers. *Journal of Political Economy*, 265–294.
- HOLTZ-EAKIN, DOUGLAS, JOULFAIAN, DAVID, & ROSEN, HARVEY S. 1994. Entrepreneurial decisions and liquidity constraints. *RAND Journal of Economics*, **25**(2), pp. 334–347.
- HOWELL, SABRINA T. 2017. Financing Innovation: Evidence from R&D Grants. *The American economic review*, **107**(4), 1136–1164.
- HOWELL, SABRINA T. 2020. Reducing information frictions in venture capital: The role of new venture competitions. *Journal of financial economics*, **136**(3), 676–694.
- HURST, ERIK, & LUSARDI, ANNAMARIA. 2004. Liquidity constraints, household wealth, and entrepreneurship. *Journal of Political Economy*, **112**(2), 319–347.
- KIHLSTROM, RICHARD E., & LAFFONT, JEAN-JACQUES. 1979. A General Equilibrium Entrepreneurial Theory of Firm Formation Based on Risk Aversion. *The Journal of political economy*, **87**(4), 719–748.
- KLETTE, TOR JAKOB, & KORTUM, SAMUEL. 2004. Innovating firms and aggregate innovation. *Journal of Political Economy*, **112**(5), 986–1018.
- KNIGHT, FRANK H. 1921. *Risk, Uncertainty and Profit*. Houghton Mifflin.
- LAZEAR, EDWARD P. 2004. Balanced skills and entrepreneurship. *American Economic Review, P&P*, **94**(2), 208–211.
- LEMLEY, MARK A. 2011. The myth of the sole inventor. *Michigan Law Review*, **110**, 709.
- LERNER, JOSH, & NANDA, RAMANA. 2020. Venture Capital’s Role in Financing Innovation: What We Know and How Much We Still Need to Learn. *The Journal of economic perspectives*, **34**(3), 237–261.
- LUCAS, ROBERT E. 1978. On the Size Distribution of Business Firms. *Bell Journal of Economics*, **9**(2), 508–523.
- MANKIW, N GREGORY, & WHINSTON, MICHAEL D. 1986. Free entry and social inefficiency. *The RAND Journal of Economics*, 48–58.
- MANNING, BENJAMIN S, ZHU, KEHANG, & HORTON, JOHN J. 2024. Automated Social Science: Language Models as Scientist and Subjects. *arXiv.org*.
- MARX, MATT, & HSU, DAVID H. 2022. Revisiting the Entrepreneurial Commercialization of Academic Science: Evidence from “Twin” Discoveries. *Management science*, **68**(2), 1330–1352.
- MERTON, ROBERT K. 1957. Priorities in scientific discovery: A chapter in the sociology of science. *American Sociological Review*, **22**(6), 635–659.
- PELLEGRINO, BRUNO. 2025. Product Differentiation and Oligopoly: A Network Approach. *The American economic review*, **115**(4), 1170–1225.

- ROBERTS, HAL, BHARGAVA, RAHUL, VALIUKAS, LINAS, JEN, DENNIS, MALIK, MOMIN M, BISHOP, CINDY, NDULUE, EMILY, DAVE, AASHKA, CLARK, JUSTIN, ETLING, BRUCE, FARIS, ROB, SHAH, ANUSHKA, RUBINOVITZ, JASMIN, HOPE, ALEXIS, D'IGNAZIO, CATHERINE, BERMEJO, FERNANDO, BENKLER, YOCHAI, & ZUCKERMAN, ETHAN. 2021. Media Cloud: Massive Open Source Collection of Global News on the Open Web. *arXiv.org*.
- ROY, ANDREW D. 1951. Some thoughts on the distribution of earnings. *Oxford Economic Papers*, **3**(2), 135–146.
- SARIRI, AMIR. 2025. The Economics of Advice: Evidence from Start-up Mentoring. *Management science*, **71**(12), 10022–10046.
- TRAAG, V. A., WALTMAN, L., & VAN ECK, N. J. 2019. From Louvain to Leiden: guaranteeing well-connected communities. *Scientific reports*, **9**(1), 5233–.
- ÅSTEBRO, THOMAS, HERZ, HOLGER, NANDA, RAMANA, & WEBER, ROBERTO A. 2014. Seeking the Roots of Entrepreneurship: Insights from Behavioral Economics. *The Journal of economic perspectives*, **28**(3), 49–69.

# Supplemental Appendix

The Online Appendix consists of the following parts.

## Appendix A Theory Extensions

This appendix collects comparative statics and auxiliary results that are not reported in the main text. The expressions are written in elasticities, which make clear how the elasticity of labor demand and the elasticity of startup success with respect to entry shape the response to shocks. See [Table 1](#) in the main text for notation.

### A.1 Idea Supply and Labor Demand Elasticities

The marginal entrepreneur may pursue a previously un-pursued idea or crowd an existing one. Entrepreneurs that pursue already-pursued ideas only crowd out incumbents, leading to no net increase in successful startups and hence no increase in labor demand. As the labor demand curve is  $D = q_0 E d(w)$ , the number of entrepreneurs enters linearly. However,  $E$  also affects  $q_0$ , which is lowered when more entrepreneurs enter the system. As such, the elasticity of labor demand with respect to the number of entrepreneurs is

$$\eta_E^D = 1 - |\eta_E^{q_0}|. \quad (12)$$

If  $\eta_E^{q_0} = -1$ , increasing the number of entrepreneurs would have no effect on labor demand. If  $\eta_E^{q_0} = 0$ , every marginal entrepreneur pursues a previously un-pursued idea and  $\eta_E^D = 1$ . *Proof.* See [Appendix B.4](#).

A related object is the wage elasticity of firm-level labor demand. As per-firm labor demand is  $d(w)$ , its elasticity with respect to the wage, in magnitude, is

$$|\eta_w^D| := -\frac{d \ln d(w)}{d \ln w} > 0. \quad (13)$$

Equivalently,  $|\eta_w^D|$  is the partial elasticity of aggregate demand  $D = q_0 E \cdot d(w)$  with respect to  $w$ , holding  $q_0$  and  $E$  fixed. Strict concavity of  $\phi$  ensures  $|\eta_w^D| > 0$ .

Another way that the startup success probability can change is through a change in the supply of ideas,  $d\kappa > 0$ , holding the number of entrepreneurs fixed. The elasticity of startup success probability with respect to  $\kappa$  has the same magnitude as the elasticity with respect to entrepreneurs, but opposite signs:

$$\eta_\kappa^{q_0} = -\eta_E^{q_0}. \quad (14)$$

This equivalence is useful when deriving comparative statics related to  $\kappa$ . *Proof.* See [Appendix B.2](#).

## A.2 Comparative Statics: Startup Costs

Consider an increase in startup costs,  $dC > 0$ . This could be due to, say, a more difficult market for startup funding. The most common source of changes is likely technological innovation that would tend to lower startup costs.

When startup costs increase, before engineers can adjust occupations, the entrepreneurship return curve,  $y(g)$ , is shifted down by  $dC$ . Employees are unaffected, and so the  $w(g)$  curve does not move. In terms of the notation for pre-adjustment changes,  $dw|g_0 = 0$  and  $dy|g_0 = -dC$ , and so the point elasticity of the entrepreneurial fraction with respect to startup costs is

$$\eta_C^g = -\frac{C/w}{\eta_g^w - \eta_g^y}. \quad (15)$$

As  $\eta_g^w - \eta_g^y > 0$ , a rise in startup costs causes engineers to flow from entrepreneurship into employment, i.e.,  $\eta_C^g < 0$ . When  $\eta_g^w - \eta_g^y$  is large, a relatively small flow of entrepreneurs is sufficient to re-establish an equilibrium, and vice versa when  $\eta_g^w - \eta_g^y$  is small.

The equilibrium change in wages with respect to a change in  $C$  is

$$\frac{\partial w}{\partial C} = g \left( -1 + \frac{1}{1-e} \eta_E^{g_0} \eta_C^g \right), \quad (16)$$

which we can re-write as

$$\frac{\partial w}{\partial C} = g \left( -1 + \frac{1}{e} \frac{|\eta_E^{g_0}|}{\eta_g^w - \eta_g^y} \right). \quad (17)$$

When  $g$  is relatively high, meaning lots of engineers are entrepreneurs, changes in  $C$  have a large effect on  $w$ . When there are relatively few entrepreneurs but many employees (low  $g$ ), higher startup costs have a smaller effect on wages since there is less total increase in startup costs to be borne by engineers as a group.

Suppose that the startup success probability was completely inelastic,  $|\eta_E^{g_0}| = 0$ . The increase in startup costs drives engineers from entrepreneurship, but because the startup success probability does not change, there is no compensating increase in success probability that would occur if  $|\eta_E^{g_0}| > 0$ . As such, a larger flow out of entrepreneurship is needed to re-establish the equilibrium, which means that employees see a larger fall in wages. With a highly elastic success probability, a smaller number of exiting engineers is needed to establish a new equilibrium, and so there is less downward wage pressure and so less pass through of startup costs.

Proposition 7 summarizes the results already discussed and contains predictions about the other outcomes of interest following a change in startup costs.

**Proposition 7** *An increase in startup costs: (1) lowers the wages of engineers, (2) lowers the retained equity of entrepreneurs, (3) raises the startup probability of success, (4) raises expected profits, (5) raises realized profits, and (6) reduces the fraction of engineers pursuing entrepreneurship.*

*Proof.* See Appendix B.6.

### A.3 Comparative Statics: Supply of Engineers

Consider an increase in the supply of engineers,  $dS > 0$ . Assume that these  $dS$  engineers initially split into employment and entrepreneurship in the same proportion as existing engineers:  $(1 - g)dS$  join the labor market, while  $gdS$  become entrepreneurs.

The elasticity of the engineer wage with respect to the supply shock at the original entrepreneurial fraction is

$$\eta_S^{w|g_0} = -\frac{\eta_E^{q_0}}{\eta_w^D}. \quad (18)$$

Despite there being more entrepreneurs, which increases demand, wages fall with a supply shock, as  $\eta_S^w|_{dg=0} < 0$ .

The elasticity of entrepreneur returns, before occupational adjustment, is

$$\eta_S^{y|g_0} = -q_0 l^* \eta_S^{w|g_0} + \frac{1}{e} \eta_E^{q_0}. \quad (19)$$

After the initial shock, engineers flow to the relatively more advantaged occupation. The elasticity of the entrepreneurial fraction with respect to  $S$  is

$$\eta_S^g = \left( \frac{1}{g\eta_w^D} + \frac{1}{e} \right) \frac{\eta_E^{q_0}}{\eta_g^w - \eta_g^y}.$$

The sign of  $\eta_S^g$  tells us whether new entrants are biased towards entrepreneurship or employment. As  $\eta_E^{q_0} < 0$ , the shock is biased towards entrepreneurship when  $|\eta_w^D| < e/g$ . If labor demand is highly elastic, when  $S$  goes up, these new engineers can be absorbed in the labor market with little reduction in wages, and so the supply shock tends to be employment-biased. The reverse is true if labor demand is highly inelastic, as wages quickly fall with more employees.

The elasticity of engineer wages to a supply shock is

$$\eta_S^w = \left( \frac{g}{e} \right) \eta_E^{q_0} (1 + \eta_S^g). \quad (20)$$

As  $\eta_E^{q_0} < 0$ ,  $\eta_S^w < 0$  if  $1 + \eta_S^g > 0$ , which is the case: even if all new engineers entered employment, the elasticity is  $\lim_{dS \rightarrow 0} \frac{\frac{E}{S+dS} - \frac{E}{S}}{(dS)/S} = -g$  and  $g < 1$ . Although a supply shock always lowers engineer wages, the amount depends in part on the elasticity of startup success probability with respect to entrepreneurs. When  $\eta_E^{q_0} = 0$ , engineer wages do not fall at all with more supply, as  $\eta_S^w = 0$ . The reason is that in the absence of idea competition, the returns to entrepreneurship do not decrease with more entrepreneurs, and so the entering engineers do as well as incumbent engineers did, pre-shock. When the startup success probability is inelastic, adding more engineer labor supply is a nearly free lunch; new entrants do not lower the wages of incumbent engineers or hurt incumbent entrepreneurs, but they do increase the number of successful startups.

Proposition 8 summarizes the other comparative static results for a positive supply shock of engineers.

**Proposition 8** *An increase in the supply of engineers: (1) lowers the wages of engineers, (2) lowers the retained equity of entrepreneurs, (3) lowers expected profits, (4) raises realized profits, (5) lowers the startup probability of success, and (6) has an ambiguous effect on entrepreneurship.*

*Proof.* See Appendix B.8.

## A.4 Comparative Statics: Supply of Ideas

Consider a positive shock to the supply of ideas,  $d\kappa > 0$ , before engineers can adjust occupations. This increases the “space” of ideas for would-be entrepreneurs, which reduces crowding and thus raises the startup success probability. This would seemingly increase the appeal of entrepreneurship, but with reduced crowding on ideas, there are more successful startups and hence increased demand for labor. As such, the relative increase in returns is what matters for predicting the flow of engineers post-shock, as was the case with the supply shock.

The effect of more ideas on wages at the original entrepreneurial fraction is

$$\begin{aligned}\eta_{\kappa}^{w|g_0} &= -\eta_{\kappa}^{q_0|g_0} / \eta_w^D \\ &= \eta_E^{q_0} / \eta_w^D.\end{aligned}$$

The effect on the returns to entrepreneurship is

$$\begin{aligned}\eta_{\kappa}^{y|g_0} &= -q_0 l^* \eta_{\kappa}^{w|g_0} + \frac{1}{e} \eta_{\kappa}^{q_0|g_0} \\ &= -q_0 l^* \eta_{\kappa}^{w|g_0} - \frac{1}{e} \eta_E^{q_0}.\end{aligned}$$

After engineers adjust occupations, the net effect on the entrepreneurial fraction is

$$\eta_{\kappa}^g = - \left( \frac{1}{g\eta_w^D} + \frac{1}{e} \right) \frac{\eta_E^{q_0}}{\eta_w^g - \eta_g^y}.$$

Whether entrepreneurship increases or decreases depends on whether  $|\eta_w^D| < e/g$ , just as in the case of a supply shock. The difference is that with the idea shock, both wage and entrepreneurship returns curves are shifting up. If labor demand is highly elastic, when  $\kappa$  goes up, the additional startups that result do not raise wages very much, and so the shock is biased towards entrepreneurship. The reverse is true if labor demand is highly inelastic, as the additional startups have a large positive effect on wages.

The net effect on wages from a positive ideas shock, after entrepreneurs adjust, is

$$\eta_{\kappa}^w = - \left( \frac{g}{e} \right) \eta_E^{q_0} (1 - \eta_{\kappa}^g). \quad (21)$$

As  $\eta_E^{q_0} < 0$ ,  $\eta_{\kappa}^w > 0$ , and so wages rise following a positive idea shock. Note that the larger  $|\eta_E^{q_0}|$ , the bigger the effect on wages. This simply reflects the fact that when startup success probability is highly elastic with respect to  $E$ , it is also highly elastic with respect to  $\kappa$ , as

these two elasticities have the same magnitude:  $\eta_\kappa^g = -\eta_S^g$  and  $\eta_\kappa^w = -\eta_S^w$ , mirroring the earlier results about the equivalence of changes in  $E$  and changes in  $\kappa$  on the startup success probability. When there is lots of idea competition, increasing  $\kappa$  is particularly welcome. Interestingly, if  $\eta_E^{g_0} = 0$ , then creating more ideas would not increase wages for the simple reason that the marginal engineer pursues ideas that already were previously unpursued and so adding more unpursued ideas does nothing.

For all the other shocks considered so far, the startup success probability was only affected through the channel of engineers entering or exiting entrepreneurship. However, in the case of the supply of ideas, the idea supply shock has a separate, independent effect on success probability. The net effect is

$$\frac{\partial q_0}{\partial \kappa} = \frac{\partial q_0}{\partial g} \frac{\partial g}{\partial \kappa} + \frac{\partial q_0}{\partial \kappa} \Big|_{g_0},$$

or in terms of elasticities,  $\eta_\kappa^{q_0} = |\eta_E^{q_0}|(1 - \eta_\kappa^g)$ .

Although wages always rise when  $\kappa$  increases at the original  $g_0$ , it is possible that the returns to entrepreneurship fall before the adjustment, i.e.,  $\eta_\kappa^{y|g_0} < 0$ . This can occur when labor costs are high, and those costs are largely passed through to entrepreneurs. For example, if  $e$  is close to 1,  $q_0$  and  $l^*$  are large, and demand is very inelastic, an increase in demand from a higher  $\kappa$  could cause a rise in wages high enough to lower the returns to entrepreneurship. However, when engineers can adjust, both  $w$  and  $y$  have to rise. The reason is that in a scenario where  $\eta_\kappa^{y|g_0} < 0$ , entrepreneurs would flow out of entrepreneurship. This would make  $\eta_\kappa^g < 0$ , and from Equation 21, wages would have to increase (and hence  $y$  would be higher in equilibrium as well).

Proposition 9 summarizes the comparative static results from a positive shock to  $\kappa$ .

**Proposition 9** *An increase in the stock of startup ideas: (1) raises the wages of engineers, (2) raises the retained equity of entrepreneurs, (3) raises expected profits, (4) lowers realized profits, (5) raises the startup probability of success, and (6) has an ambiguous effect on entrepreneurship.*

*Proof.* See Appendix B.9.

## A.5 Comparative Statics: Product Market Size

Consider an increase in the extent of the product market,  $dR > 0$ . Before engineers can change occupations, the effect on the wages of employees is  $\frac{\partial w}{\partial R} \Big|_{g_0} = \phi'(l^*) = w/R$ . As such,  $\eta_R^{w|g_0} = 1$ . For the successful entrepreneur, the effect on profits is  $\frac{\partial \pi}{\partial R} \Big|_{g_0} = \pi/R$ , and so  $\eta_R^{\pi|g_0} = 1$ . However, what matters to entrepreneurs is the effect on expected profits, which depends on the effect on  $e$  and  $q_0$ . Without any occupational adjustment,  $q_0$  stays the same. The case of equity is different: as VCs only make a return on the startup costs, their share of equity falls, as profits are larger. The entrepreneur's equity increases by  $\frac{\partial e}{\partial R} \Big|_{g_0} = (1 - e)/R$ . The overall effect on the entrepreneur's return from an expansion of the product market is  $\eta_R^{y|g_0} = 1/e$ .

With occupational adjustment, both  $\eta_R^{w|g_0}$  and  $\eta_R^{y|g_0}$  are positive, and so wages rise and engineers are shifted towards entrepreneurship. The elasticity of the entrepreneurial fraction

with respect to the extent of the product market is

$$\eta_R^g = \frac{(1-e)/e}{\eta_g^w - \eta_g^y}.$$

As  $e > 0$  and  $\eta_g^w - \eta_g^y > 0$ ,  $\eta_R^g > 0$ .

An increase in the extent of the product market draws engineers into entrepreneurship. Interestingly, this elasticity is the same magnitude (and opposite sign) as that from Equation 15, which gave the expression for  $\eta_C^g$ , or the elasticity of the entrepreneurial fraction with respect to the startup cost (note that  $(1-e)/e = C/w$ ). In other words, a 10% expansion in revenue available in the product market has the same effect as a 10% reduction in the costs of doing a startup.

The expression for the effect on the wages of engineers following adjustment is

$$\eta_R^w = 1 + \eta_g^w \eta_R^g. \quad (22)$$

As  $\eta_g^w > 0$  and  $\eta_R^g > 0$ , the elasticity of wages with respect to the extent of the product is not only positive, but it also has an elasticity greater than 1. When the product market expands, employees get a percentage wage increase the same size as the product market expansion. But then as the shock draws engineers into entrepreneurship, those employees that stay also benefit from the tighter labor market, receiving even higher wages. Engineers as a whole end up taking a larger share of the revenue compared to VCs, who still simply just get their return  $r$  on  $c$ .

Proposition 10 summarizes the comparative static results from a positive shock to  $R$ .<sup>1</sup>

**Proposition 10** *An increase in the size of the product market: (1) raises the wages of engineers, (2) raises the retained equity of entrepreneurs, (3) lowers the startup probability of success, (4) raises expected profits, (5) raises realized profits, and (6) increases the fraction of engineers pursuing entrepreneurship.*

*Proof.* See Appendix B.10.

## A.6 Comparative Statics: AI Productivity Shock

Consider a labor-augmenting productivity shock,  $d\alpha > 0$ , that replaces  $\phi(l)$  with  $\phi(\alpha l)$ , holding all other primitives fixed. The firm's first-order condition  $\alpha\phi'(\alpha l^*)R = w$  implies that efficiency units  $u^* \equiv \alpha l^*$  depend only on the efficiency wage  $w/\alpha$ , with  $u^* = \psi(w/\alpha)$  and  $\psi \equiv (\phi')^{-1}(\cdot/R)$ . Per-firm human labor is  $l^* = \psi(w/\alpha)/\alpha$ , and the profit function inherits the same dependence,  $\pi = \Pi(w/\alpha)$ , with  $\Pi'(x) = -\psi(x)$  by the envelope theorem.

Log-differentiating the occupational indifference condition  $w = q_0\Pi(w/\alpha) - C$ , using

---

<sup>1</sup>When the product market expands, startups become *smaller*, as engineer employees became expensive. While this perhaps seems counterintuitive, the phenomenon of startups with very high valuations but small headcounts is much remarked upon in Silicon Valley. For example, when WhatsApp was acquired by Facebook for \$14 billion, it had 42 employees. Presumably, this remarkable outcome is related to the productivity of elite engineers, but it also suggests WhatsApp was economizing on expensive engineer labor.

$(1 - g)/g = q_0 l^*$  and  $q_0 \pi/w = 1/e$ , yields

$$d \ln w = \frac{g}{e} d \ln q_0 + (1 - g) d \ln \alpha. \quad (23)$$

The first term recovers the wage–success-probability tradeoff familiar from Proposition 8; the second captures the direct labor-market impact of the productivity gain.

Market clearing  $S = E(1 + q_0 l^*)$ , with  $E = \kappa F(q_0)$  where  $F(q_0) \equiv \int_{q_0}^1 (q/q_0) f(q) dq$  and  $d \ln l^* = -|\eta_w^D| d \ln w + (|\eta_w^D| - 1) d \ln \alpha$ , log-differentiates to

$$(g + h) d \ln q_0 = -(1 - g)|\eta_w^D| d \ln w + (1 - g)(|\eta_w^D| - 1) d \ln \alpha, \quad (24)$$

where  $h \equiv \kappa q_0 f(q_0)/E$  denotes the density of ideas at the margin relative to the mass of entrepreneurs. Substituting (23) into (24) and solving gives the elasticity of the startup success probability with respect to the shock,

$$\eta_\alpha^{q_0} = \frac{(1 - g)(|\eta_w^D| g - 1)}{D}, \quad D \equiv g + h + (1 - g) \frac{|\eta_w^D| g}{e} > 0. \quad (25)$$

The sign of  $\eta_\alpha^{q_0}$  equals the sign of  $|\eta_w^D| g - 1$ .

The remaining elasticities follow. Using  $d \ln g = -(1 - g)(d \ln q_0 + d \ln l^*)$ ,

$$\eta_\alpha^g = \frac{(1 - g)(1 - |\eta_w^D| g)(1 + h)}{D}, \quad (26)$$

with sign opposite to  $\eta_\alpha^{q_0}$ . Engineers per successful firm satisfy  $\eta_\alpha^{l^*} = (|\eta_w^D| g - 1)(g + h)/D$ , the same sign as  $\eta_\alpha^{q_0}$ . The wage elasticity is

$$\eta_\alpha^w = (1 - g) + \frac{g}{e}(1 - g) \frac{|\eta_w^D| g - 1}{D}. \quad (27)$$

Factoring yields

$$\eta_\alpha^w = \frac{(1 - g)}{e D} \left[ eh + g(|\eta_w^D| - (1 - e)) \right], \quad (28)$$

so  $\eta_\alpha^w > 0$  if and only if  $g|\eta_w^D| + eh > g(1 - e)$ . This condition is satisfied whenever  $|\eta_w^D| \geq 1 - e$ , and a fortiori whenever  $|\eta_w^D| g \geq 1$ . At  $|\eta_w^D| g = 1$ ,  $\eta_\alpha^w = 1 - g$  exactly; for  $|\eta_w^D| g < 1$  the elasticity lies below  $1 - g$  (when positive); for  $|\eta_w^D| g > 1$  it exceeds  $1 - g$ .

Expected profits satisfy  $q_0 \pi = w + C$ , so  $\eta_\alpha^{q_0 \pi} = (w/(q_0 \pi)) \eta_\alpha^w = e \cdot \eta_\alpha^w$ . Retained equity  $e = 1 - C/(q_0 \pi)$  has  $\eta_\alpha^e = (1 - e) \eta_\alpha^w$ . Both inherit the wage sign. Realized profits  $\pi = (w + C)/q_0$  rise unconditionally:

$$\eta_\alpha^\pi = (1 - g) \left[ e - \frac{(1 - g)(|\eta_w^D| g - 1)}{D} \right]. \quad (29)$$

For  $|\eta_w^D| g \leq 1$ , the bracket exceeds  $e$  (both terms non-negative). For  $|\eta_w^D| g > 1$ ,  $eD - (1 - g)(|\eta_w^D| g - 1) = eg + eh + (1 - g) > 0$ , so the bracket is again positive. Realized profits rise regardless of parameter region—the one unambiguous labor-market-side effect of the shock.

In the limit  $\alpha \rightarrow \infty$ , the efficiency wage  $w/\alpha \rightarrow 0$ , so  $u^* \rightarrow \psi(0) = S$  is finite while  $l^* = u^*/\alpha \rightarrow 0$ . Market clearing  $S = E(1 + q_0 l^*)$  then forces  $E \rightarrow S$  and, via  $E = \kappa F(q_0)$ ,  $q_0 \rightarrow F^{-1}(S/\kappa) > 0$ ; the idea threshold settles at a strictly positive limit, and  $g = E/S \rightarrow 1$ . This is the “entrepreneurship-only” limit of Corollary 6.

## Appendix B Proofs

### B.1 Proof of Proposition 1

**Proof.** The plan of the proof is to (1) show that the returns to employment and entrepreneurship move in opposite directions with respect to  $g$ , the fraction of engineers choosing entrepreneurship, and (2) show that as nearly everyone pursues employment, or  $g \rightarrow 0$ , the returns to entrepreneurship are greater than the returns to employment, and vice-versa when nearly everyone pursues entrepreneurship,  $g \rightarrow 1$ . This ensures that the two return curves cross in the unit interval, i.e., that there exists some equilibrium  $g \in (0, 1)$ . With this  $g$ , the returns to employment and entrepreneurship are equalized, which is required for an equilibrium. As the return curves are monotonic, this equilibrium is unique.

**Opposite signs for the slopes of  $w$  and  $y$  in  $g$ .** A larger  $g$  means more entrepreneurs and hence more successful firms demanding labor; it also means fewer employees in the labor market. The higher the demand and the lower the supply, the higher wages are, and so  $w'(g) > 0$ . More entrepreneurs also reduces the returns to entrepreneurship. With more entrepreneurs,  $q_0$  goes down (from the entrepreneur-sorting condition that assigns  $q/q_0$  entrants to each pursued idea). Because of the higher wages for employees at this higher  $g$ , profits go down. Finally,  $e$  goes down, as lower expected profits means that entrepreneurs must give up more equity to obtain seed funding. As such,  $y'(g) < 0$ , where  $y(g) \equiv eq_0\pi$  denotes the expected return to entrepreneurship.

**Returns cross for some  $g$  in  $(0, 1)$ .** Although the slopes of  $w(g)$  and  $y(g)$  have opposite signs, for an equilibrium to exist, they must cross in the unit interval, i.e., there must be a  $g \in (0, 1)$  such that  $w(g) = y(g)$ . First, consider the case when  $g \rightarrow 0$ , meaning almost every engineer is an employee. Let  $\bar{q} \leq 1$  be the startup success probability when there are almost no entrepreneurs. By assumption,  $C/R < \bar{q}$  (recall that the proposition states that the equilibrium exists only for a sufficiently small  $C/R$ ). As the number of entrepreneurs gets arbitrarily small, the number of successful startups gets arbitrarily small, and thus the number of workers per firm approaches  $S$ . The returns to entrepreneurship minus the returns to employment is

$$\lim_{l \rightarrow S} \bar{q} (R\phi(l) - R\phi'(l)l) - C - R\phi'(l) = \bar{q}R - C.$$

Recall that by assumption, as  $l \rightarrow S$ ,  $\phi(l) \rightarrow 1$  and  $\phi'(l) \rightarrow 0$ . As  $\bar{q}R - C > 0$ , as the number of entrepreneurs gets arbitrarily small, the returns to entrepreneurship are greater than the returns to employment.

Now we consider the case when  $g \rightarrow 1$ , meaning that nearly all engineers pursue entrepreneurship. With almost no engineers available to be employees,  $l \rightarrow 0$ . The returns to entrepreneurship minus the returns to employment is

$$\lim_{l \rightarrow 0} \underline{q} (R\phi(l) - R\phi'(l)l) - C - R\phi'(l) = -C - Rp,$$

where  $\underline{q} = \lim_{g \rightarrow 1} q_0(g)$  and  $p = \lim_{l \rightarrow 0} \phi'(l)$ . As  $p > 0$ ,  $C > 0$  and  $R > 0$ , the returns to

entrepreneurship are lower than the returns to employment. We can now conclude that there is some  $g \in (0, 1)$  such that the returns to employment and entrepreneurship are equalized, or  $w(g) = y(g)$ . As  $w(\cdot)$  and  $y(\cdot)$  are monotonic, the equilibrium is unique. ■

## B.2 Derivation of the elasticity of startup success probability with respect to ideas

Using  $E = \kappa \int_{q_0}^1 \frac{q}{q_0} f(q) dq$ , following a small increase in  $\kappa$ , but keeping the number of entrepreneurs fixed, we can write

$$0 = \frac{E}{\kappa} + \left( \kappa f(q_0) + \kappa \frac{E}{q_0} \right) \frac{\partial q_0}{\partial \kappa},$$

and as an elasticity,

$$\eta_{\kappa}^{q_0} = \frac{1}{1 + \kappa f(q_0) q_0 E^{-1}}.$$

## B.3 Profits are always positive

Because  $\phi(l)$  is concave,

$$\phi(l^*) = \int_0^{l^*} \phi'(x) dx > l^* \phi'(l^*),$$

and so  $R(\phi(l^*) - \phi'(l^*)l^*) = \pi > 0$  (recall that employees are paid their marginal product).

## B.4 Derivation of Equation 12

The demand curve is  $D = E q_0 d(w)$ . Differentiating with respect to the number of entrepreneurs gives

$$\begin{aligned} \frac{\partial D}{\partial E} &= q_0 d(w) + E d(w) \frac{\partial q_0}{\partial E} \\ \frac{1}{D} \frac{\partial D}{\partial E} &= \frac{1}{E} + \frac{1}{q_0} \frac{\partial q_0}{\partial E} \\ \frac{E}{D} \frac{\partial D}{\partial E} &= 1 + \frac{E}{q_0} \frac{\partial q_0}{\partial E} \\ \eta_E^D &= 1 + \eta_E^{q_0}. \end{aligned}$$

Since  $\eta_E^{q_0} \leq 0$ , this is equivalent to  $1 - |\eta_E^{q_0}|$ .

## B.5 Proof of Lemmas

Two lemmas are useful for simplifying the comparative static prediction proofs. Lemma 11 shows that  $w$  and  $q_0$  move in opposite directions for all shocks (except those to the supply of

engineers and/or the supply of ideas). Lemma 12 relates the wages of engineers to changes in the supply of engineers and the supply of startup ideas.

**Lemma 11** *Any small change in an exogenous variable (with respect to  $S$  and  $\kappa$ ) that raises  $w$  decreases  $q_0$  and vice versa.*

**Proof.** The labor market clearing condition  $S = E + Eq_0l^*$  can be written as

$$d(w) = \frac{S}{\kappa} \frac{1}{\int_{q_0}^1 qf(q)dq} - \frac{1}{q_0}.$$

Let  $\gamma$  be some exogenous variable that does not affect  $\kappa$  or  $S$  but does affect  $w$  and  $q_0$ . If we differentiate the market clearing condition by  $\gamma$ , we have

$$d'(w) \frac{\partial w}{\partial \gamma} = \frac{\partial q_0}{\partial \gamma} \left( 1 + \frac{S}{\kappa} \frac{f(q_0)q_0^3}{\left(\int_{q_0}^1 qf(q)dq\right)^2} \right) \frac{1}{q_0^2}.$$

On the right-hand side, the last two factors are both positive. And since  $d'(w) < 0$ ,  $\frac{\partial w}{\partial \gamma}$  and  $\frac{\partial q_0}{\partial \gamma}$  have opposite signs. ■

**Lemma 12** *Engineer wages are decreasing the ratio of engineers to ideas.*

**Proof.** Let the ratio of engineers to startup ideas be  $a = S/\kappa$ . Using the fact that  $g = \frac{1}{1+q_0d(w)}$ , the market clearing condition can be written as

$$d(w) = a \frac{1}{\int_{q_0}^1 qf(q)dq} - \frac{1}{q_0}.$$

Differentiating with respect to  $a$ , we have

$$d'(w) \frac{\partial w}{\partial a} = \frac{1}{\int_{q_0}^1 qf(q)dq} + \frac{\partial q_0}{\partial a} \left( \frac{1}{q_0^2} + \frac{af(q_0)}{\left(\int_{q_0}^1 qf(q)dq\right)^2} \right).$$

Assume that  $\frac{\partial w}{\partial a} > 0$ . Since  $d'(w) < 0$ , the entire right-hand side must be negative, which implies that  $\frac{\partial q_0}{\partial a} < 0$ . However, if we differentiate the expression for employee wages by  $a$ , we get

$$\frac{\partial w}{\partial a} (1 + d(w)q_0) = \pi \frac{\partial q_0}{\partial a}, \tag{30}$$

which implies that  $\frac{\partial w}{\partial a}$  and  $\frac{\partial q_0}{\partial a}$  have the same sign, contrary to our original assumption. Therefore,  $\frac{\partial w}{\partial a} < 0$ . ■

## B.6 Proof of Proposition 7

**Proof.** For (1), starting from  $C = q_0\pi - w$ , and differentiating by  $C$ , we have

$$1 = \pi \frac{\partial q_0}{\partial C} - [1 + q_0 d(w)] \frac{\partial w}{\partial C}. \quad (31)$$

Assume that  $\frac{\partial q_0}{\partial C} < 0$ . For Equation 31 to hold, it would imply that  $\frac{\partial w}{\partial C} < 0$ , but this is a contradiction by Lemma 11, and therefore  $\frac{\partial q_0}{\partial C} > 0$ , which is claim (3). This implies that  $\frac{\partial w}{\partial C} < 0$  by Lemma 11, which is claim (1). For claim (2), as  $e = w/(w + C)$ , we know that retained equity falls when startup costs increase. For claim (4), since  $q_0$  goes up and  $w$  goes down, expected profits,  $q_0\pi$ , go up. For claim (5), as  $w$  is lower, realized profits go up as well. For claim (6), since  $w$  goes down,  $d(w)$  goes up. As we know from claim (3),  $q_0$  goes up, and therefore the entrepreneurial fraction goes down. ■

## B.7 Derivation of supply shock effects

Before any adjustment in occupations, market clearing in the labor market requires that

$$\begin{aligned} dw &= \frac{-(1-g)dS + \frac{\partial D}{\partial E}gdS}{|\partial D/\partial w|} \\ \frac{S}{w} \frac{dw}{dS} &= \frac{-S(1-g) + \frac{\partial D}{\partial E}gS}{w|\partial D/\partial w|} \\ \eta_S^w &= \frac{-1 + \eta_E^D}{|\eta_w^D|} \\ \eta_S^w &= \frac{-1 + 1 + \eta_E^{q_0}}{|\eta_w^D|} \\ \eta_S^w &= \frac{\eta_E^{q_0}}{|\eta_w^D|}. \end{aligned}$$

## B.8 Proof of Proposition 8

**Proof.** The lowering of wages, which is claim (1), follows from Lemma 12, as  $\frac{\partial a}{\partial S} > 0$  (recall that  $a = S/\kappa$ ). For claim (2), from the fact that  $e = w/(w + C)$ , we know that retained equity falls as well since  $\frac{\partial}{\partial w} \left[ \frac{w}{w+C} \right] > 0$ . For claim (3), from the occupational indifference condition,  $\frac{\partial w}{\partial S} = \frac{\partial}{\partial S} q_0\pi < 0$  (making use of the fact that wages fall). For claim (4),  $\frac{\partial \pi}{\partial S} = -d(w) \frac{\partial w}{\partial S} > 0$ , as  $d(w) > 0$  (and again making use of (1)). For claim (5), from the occupational indifference condition, we know that

$$\frac{\partial w}{\partial S} = \frac{\partial}{\partial S} [q_0\pi],$$

and from the envelope theorem,  $\frac{\partial \pi}{\partial S} = -d(w) \frac{\partial w}{\partial S}$ , and so

$$\frac{\partial w}{\partial S} (1 + q_0 d(w)) = \pi \frac{\partial q_0}{\partial S},$$

and since  $\frac{\partial w}{\partial S} < 0$ ,  $\frac{\partial q_0}{\partial S} < 0$ . For claim (6), as  $w$  goes down,  $d(w)$  goes up, but as  $q_0$  goes down, the overall effect on the entrepreneurship fraction,  $g$ , is ambiguous. ■

## B.9 Proof of Proposition 9

**Proof.** For claim (1), when  $\kappa$  increases, the ratio of engineers to ideas decreases ( $a = S/\kappa$ ), and so it directly follows from Lemma 12 that engineer wages increase. For claim (2), as startup costs have not changed, because  $e = w/(w + C)$ , it follows that retained equity increases. For claim (3), expected profits rise since wages rise, i.e.,  $\frac{\partial w}{\partial \kappa} = \frac{\partial}{\partial \kappa}[q_0\pi] > 0$ . However, for claim (4), realized profits fall because wages increase. For claim (5), recall that

$$\frac{\partial w}{\partial a}(1 + d(w)q_0) = \pi \frac{\partial q_0}{\partial a}, \quad (32)$$

which implies that  $\frac{\partial w}{\partial a}$  and  $\frac{\partial q_0}{\partial a}$  have the same sign. From claim (1), we know that  $\frac{\partial w}{\partial a} < 0$ , and hence that  $\frac{\partial q_0}{\partial a} < 0$ . Therefore  $\frac{\partial q_0}{\partial \kappa} > 0$ , meaning that the startup probability of success increases. For claim (6), because  $w$  goes up,  $d(w)$  goes down, but  $q_0$  goes up. As such, the effect of a change in the stock of ideas on  $q_0d(w)$  is ambiguous, and hence the effect on entrepreneurship is ambiguous. ■

## B.10 Proof of Proposition 10

**Proof.** For claim (1), differentiating the expression for engineer wages by  $R$ , we have

$$\frac{\partial w}{\partial R} = q_0\phi(l) + \pi \frac{\partial q_0}{\partial R}. \quad (33)$$

Assume that  $\frac{\partial w}{\partial R} < 0$ . As  $q_0\phi(l) > 0$ , for the equation above to hold,  $\pi \frac{\partial q_0}{\partial R} < 0$ , but by Lemma 11,  $\frac{\partial q_0}{\partial R} > 0$ , given the assumption that  $\frac{\partial w}{\partial R} < 0$ . As startup costs have not changed, claim (2) follows directly from the condition that  $e = w/(w + C)$  and claim (1). For claim (3), because wages rise,  $q_0$  falls. For claim (4) expected profits rise, as  $\frac{\partial w}{\partial R} = \frac{\partial}{\partial R}[q_0\pi]$ . And for claim (5), the effect of a larger product market on realized profits is also positive, as expected profits rose despite  $q_0$  falling. For claim (6), as both  $q_0$  and  $d(w)$  go down,  $g$  goes up. ■

## B.11 Proof of Proposition 5

**Proof.** Effects (1)–(3) follow from (28):  $\text{sign}(\eta_\alpha^w) = \text{sign}(\eta_\alpha^e) = \text{sign}(\eta_\alpha^{q_0\pi}) = \text{sign}(g|\eta_w^D| + eh - g(1 - e))$ . Effect (4) follows from (29) and the positivity argument that accompanies it in Appendix A.6. Effects (5) and (6) follow from  $\text{sign}(\eta_\alpha^{q_0}) = \text{sign}(|\eta_w^D|g - 1)$  and  $\text{sign}(\eta_\alpha^g) = \text{sign}(1 - |\eta_w^D|g)$ , which can be read off from (25) and the expression for  $\eta_\alpha^g$  in Appendix A.6. ■

## B.12 Proof of Proposition 2

**Proof.** The net social benefit (total revenue minus startup costs) from the decentralized allocation of engineers is

$$\begin{aligned}
&= q_0 g S[R\phi(l^*)] - g S C \\
&= g S (q_0 [R\phi(l^*)] - C) \\
&= g S (q_0 [\pi + w l^*]) - C \\
&= g S (q_0 [\pi + w l^*]) - \pi q_0 + w \\
&= g S (w [1 + q_0 l^*]) \\
&= g S \left( \frac{w}{g} \right) \\
&= S w.
\end{aligned}$$

■

## B.13 Proof of Proposition 3

**Proof.** For ease of exposition, let  $dE = 1$ , i.e., a single engineer moves to entrepreneurship. This has a direct cost of  $C$ . If the marginal engineer does happen to choose a “new” idea to pursue and it is successful, then there is  $R\phi(l)$  more output. However, producing this output costs  $w l^*$ , as engineers must be drawn from other successful firms. The marginal engineer offers the same social benefit in entrepreneurship as in employment if

$$\begin{aligned}
(1 - |\eta_E^{q_0}|) (q_0 \pi) - C &= w, \\
(1 - |\eta_E^{q_0}|) (q_0 \pi) &= w + C, \\
(1 - |\eta_E^{q_0}|) (q_0 \pi) &= q_0 \pi, \\
1 - |\eta_E^{q_0}| &= 1,
\end{aligned}$$

or when the startup success probability with respect to the number of entrepreneurs is completely inelastic,  $|\eta_E^{q_0}| = 0$ . ■

## B.14 Proof of Proposition 4

**Proof.** By assumption,  $n^*(q_0) = n(q_0)$ , and since  $p < 1$ ,  $n^{*'}(q_0) > n'(q_0)$ . There exists a  $q'$  such that for all  $q \in (q_0, q')$ ,  $n^*(q) > n(q)$ . As the area under  $n(q)$  must equal the area under  $n^*(q)$ , the two curves must eventually cross. And because both have the same value at  $q_0$  and are monotonically increasing (but do not have the same slopes everywhere), they cross only once. As  $n(q)$  has a constant slope, and  $n^*(q)$  has a greater slope at  $q_0$ ,  $n^*(q)$  crosses from above. Let  $\hat{q}$  be that crossing point. As  $\int_{q_0}^{\hat{q}} n^*(q) f(q) dq > \int_{q_0}^{\hat{q}} n(q) f(q) dq$ , and since the area under the two curves must be equal, the optimal curve  $n^*(q)$  is everywhere below the  $n(q)$  curve for  $q > \hat{q}$ . ■

# Appendix C Data Appendix

## C.1 Clustering Implementation

This appendix provides implementation details for the clustering procedure summarized in the main text.

**Weighted Jaccard similarity:** Each firm’s keyword portfolio defines a sparse vector over the universe of distinct keyword phrases in the sample, where each entry is the normalized budget weight  $w_{ik}$  that firm  $i$  assigns to keyword  $k$ . The numerator of the weighted Jaccard index captures the extent to which two firms allocate spending to the same buyer-intent terms; the denominator normalizes by the union of their combined portfolios. A value near one implies the two firms target nearly identical customer search intents with similar spending priorities.

**Mutual  $k$ -nearest-neighbor graph:** For each firm we retain the  $k$  peers with the highest Jaccard scores, producing a directed nomination list. We symmetrize by keeping only mutual nominations and assigning each undirected edge the mean of the two directed Jaccard scores. A high  $k$  connects each firm to a broader set of peers, smoothing over noise but risking that distinct niches merge; a low  $k$  yields finer but noisier clusters.

**Leiden–CPM community detection:** The Leiden algorithm (Traag *et al.*, 2019) partitions the mutual  $k$ -NN graph by maximizing the Constant Potts Model (CPM) objective:

$$\mathcal{H}_{\text{CPM}} = \sum_c \left[ e_c - \gamma \binom{n_c}{2} \right],$$

where  $e_c$  is the sum of edge weights within community  $c$ ,  $n_c$  is the number of nodes in community  $c$ , and  $\gamma$  is the resolution parameter. The first term rewards dense within-community connectivity; the second penalizes community size. Because  $\gamma$  enters as an absolute density threshold rather than scaling with the graph, cluster granularity remains interpretable as the sample grows.

**Parameter selection and validation:** We select  $k$  and  $\gamma$  over a grid, choosing the configuration that maximizes B<sup>3</sup> F-score on the hand-labeled benchmark described in the main text. We further require that the winning configuration passes basic health checks. We require that the largest connected component must not swallow all firms, and singletons must remain a small share of the graph.

## C.2 Theory-Data Correspondence: Media Mentions as a Proxy for Idea Quality

The model predicts that an idea of quality  $q$  attracts  $n(q) = q/q_0$  entrepreneurs in equilibrium. The empirical regressions replace  $q$  with media mentions  $m$ , a proxy for perceived

idea quality. This subsection states the assumptions under which the regression coefficients recover the theoretical objects and characterizes the bias when those assumptions fail.

Let  $m = h(q)$  where  $h : [0, 1] \rightarrow \mathbb{R}_+$  is strictly increasing and differentiable. Under the model, the log-log relationship between entry and true quality is

$$\log n(q) = \log q - \log q_0,$$

which has unit elasticity. Substituting  $q = h^{-1}(m)$  gives  $\log n = \log h^{-1}(m) - \log q_0$ . The population regression of  $\log n$  on  $\log m$  recovers a slope of

$$\frac{d \log n}{d \log m} = \frac{1}{\varepsilon_h(q)}, \quad \text{where } \varepsilon_h(q) \equiv \frac{d \log m}{d \log q} = \frac{q h'(q)}{h(q)}$$

is the log-elasticity of media mentions with respect to true quality. The slope equals unity (the theoretical elasticity) if and only if  $\varepsilon_h(q) = 1$  throughout the support. When  $h$  is a power function,  $m = Aq^\alpha$ , the log-elasticity is constant at  $\varepsilon_h = \alpha$ , and the slope is  $1/\alpha$ ; it equals one only when  $\alpha = 1$ , i.e., when media mentions are proportional to true quality.

The direction of bias depends on whether  $\varepsilon_h$  is above or below one. When  $\varepsilon_h < 1$  (a one-percent increase in true quality raises media mentions by less than one percent), the measured elasticity of entry with respect to  $m$  exceeds the true elasticity with respect to  $q$ : the regression overstates how steeply entry responds to quality. When  $\varepsilon_h > 1$ , the measured elasticity understates the true response. A sufficient condition for  $\varepsilon_h \leq 1$  is that  $h$  is concave with  $h(0) \geq 0$ , since concavity implies  $h(q) \geq h(0) + q h'(q) \geq q h'(q)$ . This condition is plausible if the press under-covers ideas whose viability depends on technical details rather than consumer appeal. The reverse,  $\varepsilon_h > 1$ , would arise if media attention amplifies the top of the quality distribution, as when journalists cover “moonshot” ideas disproportionately.

For the congestion elasticity  $\eta$  estimated in the success regression, the main threat to interpretation is omitted-variable bias. If unobserved niche characteristics (low barriers, commoditized technology) simultaneously raise media salience, inflate entry, and reduce per-firm success, the OLS estimate can overstate  $|\eta|$ . We discuss this endogeneity concern in the main text and treat the full-sample Poisson estimate  $\hat{\eta} = -0.18$  as a conservative lower bound and the restricted-sample log-log estimate  $\hat{\eta} = -0.86$  as an upper bound on the degree of congestion.

## Appendix D Additional Figures and Tables

**Figure D1:** Marketing keyword prompt and agent traits used in the LLM survey

```
1 q1 = QuestionList(  
2   question_text = ""  
3   Read the following company description and then draft 10 non-branded  
4     Google Ads keyword search queries that would most effectively reach  
5     customers searching for a competing product or service. If the  
6     description is not detailed enough, infer the company's core  
7     offering as best as you can. You must strictly follow these rules  
8     in drafting the keyword search queries:  
9  
10    1) Return EXACTLY 10 search queries. Lowercase. No quotes. 1-3 words  
11     per search query.  
12    3) No locations unless the description itself is location-specific.  
13    4) Avoid generic fluff like "best" and "cheap".  
14    5) Prefer concrete buyer intent (e.g., "meal kit delivery", not "  
15     recipes").  
16    6) Singular nouns unless plural is industry-standard.  
17    7) De-duplicate; if two items would be synonyms, keep the clearer  
18     one.  
19    8) Prefer modifier-before-noun phrasing (e.g., "green car", not "car  
20     green").  
21  
22    Here is the company description: "{{ scenario.company_description }}"  
23    "" ,  
24    max_list_items = 10,  
25    min_list_items = 10,  
26    question_name = "marketing_keywords"  
27 )
```

(a) Survey question used to generate keywords

```
1 agents = AgentList(  
2   Agent(  
3     traits = {  
4       "founding_year": founding_year ,  
5       "persona": f"You are a Google AdWords specialist working in {  
6         founding_year}. Act as if the current year is {  
7         founding_year}.",  
8       "expertise": "You pick high-intent, non-branded search queries  
9         for early-stage startups."  
10    }  
11  ) for founding_year in founding_years  
12 )
```

(b) Google AdWords agent traits

Notes: **Figure D1a** shows the prompt used to extract non-branded marketing keywords a Google AdWords specialist would have bid on to reach their customers. The portion “{{ scenario.company\_description }}” is the placeholder for company description. **Figure D1b** endow the agent with the desired role as a Google AdWords specialist, where “{founding\_year}” is replaced by the company’s year of founding.

Figure D2: Technology startup classifier LLM survey

```
1 question = QuestionYesNo(  
2     question_text = ""  
3 You are evaluating whether a company qualifies as an innovation-driven,  
4     venture-fundable startup versus a traditional non-venture-fundable  
5     business. Read the company description carefully and determine if it  
6     meets the criteria for a venture-fundable startup.  
7  
8 *Decision checklist (apply in order)*  
9  
10 1) Core tech/science? The product itself is scalable technology or novel  
11     science  
12     - e.g., SaaS/app/API, digital marketplace/platform, AI/data/ML models,  
13         developer tools;  
14     - robotics/autonomy/IoT with proprietary software;  
15     - novel materials/advanced manufacturing processes;  
16     - biotech/life-science platforms (discovery, diagnostics, synbio);  
17     - fintech/insurtech infrastructure (rails, underwriting engines,  
18         issuing APIs).  
19  
20     If the core is a local service or facility (clinic, lender, realtor  
21         office, senior home, agency), this fails (answer No).  
22  
23 2) Scale intent/economics? National/global reach, low marginal cost,  
24     platform/network effects, subscriptions/APIs, IP/data moats. (Helpful  
25     cues; not required if tech/science core is explicit.)  
26  
27 3) Default rule: If the description doesn't clearly establish (1), answer  
28     No.  
29  
30 *Include archetypes (usually Yes)*  
31  
32 - Software product (SaaS/app/API), digital platform/marketplace.  
33 - Robotics/edge/IoT products with proprietary software for multi-site  
34     deployment.  
35 - New materials/chemistry/semiconductor/battery tech intended for broad  
36     B2B adoption.  
37 - Biotech platforms (e.g., screening, design, tooling) rather than a local  
38     clinic.  
39 - Fintech/insurtech rails, issuing/underwriting/claims automation  
40     platforms.  
41  
42 *Exclude archetypes (usually No)*  
43  
44 - Facility-bound/local services: dental/medical/rehab clinics, senior  
45     homes, urgent care, restaurants, gyms, salons, repair shops.  
46 - Traditional vendors: mortgage lenders/originators, realtor/brokerage  
47     offices, insurance/wealth agencies, property managers.  
48 - Agencies/consultancies/dev-shops/staffing firms; tutoring centers.  
49 - Commodity hardware or D2C retail with no proprietary software/platform/  
50     IP.  
51 - "We use AI/an app" inside a conventional service != tech-core product.
```

```

35
36 *Edge switches (decide consistently)*
37
38 - SaaS/marketplace for clinics -> Yes; the clinics themselves -> No.
39 - Telehealth platform (onboards many providers/patients) -> Yes; single
    telehealth practice -> No.
40 - Robotics product (hardware+software sold broadly) -> Yes; warehouse
    using robots to run itself -> No.
41 - Mortgage comparison/lending infrastructure -> Yes; mortgage lender
    office -> No.
42
43 *Mini-examples (description -> label)*
44
45 -"API-first payments platform issuing virtual cards for developers" -> Yes
46 -"Autonomous mobile robots with proprietary vision stack for warehouses"
    -> Yes
47 -"Novel silicon-anode battery material licensed to cell makers" -> Yes
48 -"AI platform matching patients to dentists nationwide" -> Yes
49 -"Practice-management SaaS for dental clinics" -> Yes
50 -"Local dental clinic offering Invisalign and implants" -> No
51 -"Independent mortgage lender serving our county" -> No
52 -"Boutique real-estate brokerage with top agents" -> No
53 -"Addiction treatment center with 40 beds" -> No
54 -"Marketing agency using AI for copywriting" -> No
55 -"Electronics retailer with an online store" -> No
56
57 Answer with a single word: "Yes" or "No"
58
59 Company: {{ scenario.name }}
60 Description: {{ scenario.company_description }}
61     "",
62     question_name = "is_tech_company"
63 )

```

Notes: This figure shows the prompt used to classify whether a company qualifies as a technology startup. The classifier returns a binary Yes/No label. The portions “{{ scenario.name }}” and “{{ scenario.company\_description }}” are placeholders for the company name and description, respectively.

**Table D1:** Changes in Sample Due to Sample Attrition

	N	Total funding (Million USD)	Majority state	Founding date	Employee count
1. Founded 2012–2018	826,489	4.22 (-10.49)	CA	2015-02-28 (2.37)	66.51 (0.59)
2. Primary role = company	798,295	3.92 (-10.64)	CA	2015-02-26 (1.69)	64.68 (-0.48)
3. USA only	280,693	5.81 (-9.61)	CA	2015-01-16 (-8.52)	54.34 (-6.31)
4. Homepage URL present	277,379	5.85 (-9.58)	CA	2015-01-16 (-8.57)	54.37 (-6.29)
5. Survived $\geq$ 12 months	277,039	5.86 (-9.58)	CA	2015-01-16 (-8.56)	54.41 (-6.27)
6. No Consulting/Property Mgmt	234,437	6.74 (-9.10)	CA	2015-01-18 (-8.10)	57.42 (-4.50)
7. Description non-empty and length $\geq$ 50	137,272	11.39 (-6.62)	CA	2015-01-26 (-5.79)	63.07 (-1.27)
8. Tech companies	57,171	19.51 (-2.39)	CA	2015-02-14 (-1.07)	55.26 (-4.98)
9. Non-singleton clusters	53,558	20.37 (-2.00)	CA	2015-02-15 (-0.80)	56.36 (-4.37)
10. At least 1 exit	37,818	25.27 (0.00)	CA	2015-02-19 (0.00)	65.50 (0.00)

Notes: Each row reports summary statistics after applying the corresponding filter to the Crunchbase organizations data. Employee count is computed as the midpoint of the reported range in Crunchbase (e.g., 10-50 is coded as 30) and excludes non-numeric entries (e.g., "unknown"). Majority state is the modal State code at each stage. Parentheses report t-statistics for differences in means relative to the final sample, computed using the stage-specific and final-sample variances and sample sizes (treated as independent samples).

**Table D2:** Amazon ECS Robustness Shock: Entry and Exit-Rate Effects

	Log entrants	Exit rate by 6 years
<i>Panel A: Amazon ECS (post = 2015)</i>		
Post $\times$ exposure	0.2084** (0.0918)	0.0063 (0.0163)
Observations	9,132	9,132
Pre-trend test ( $p$ -value)	0.011	

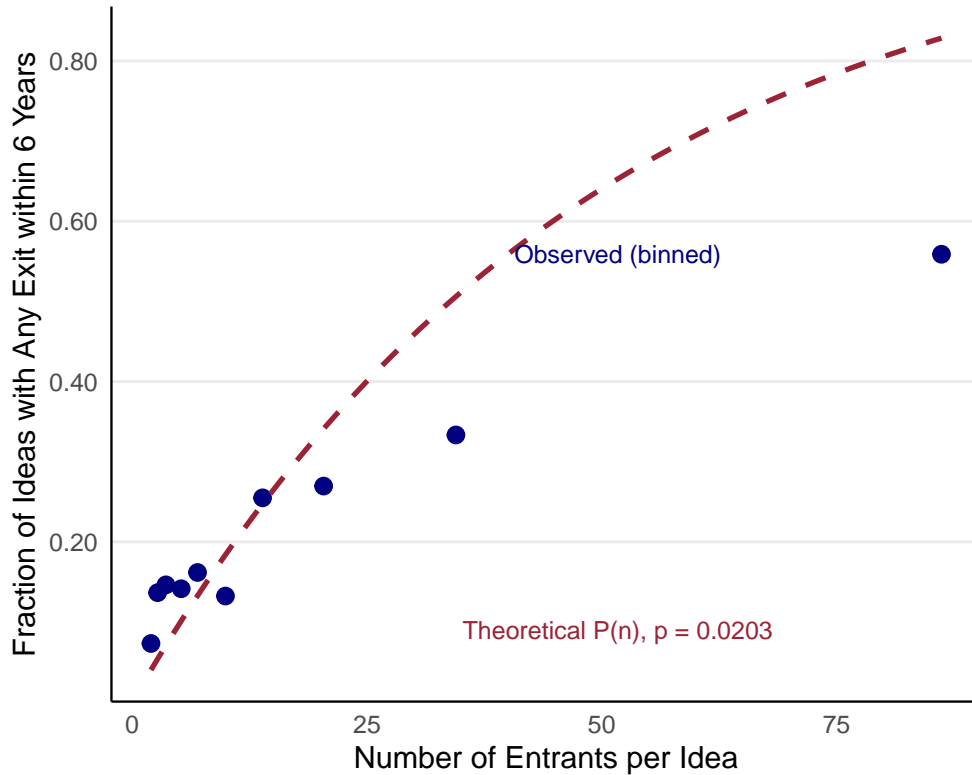
Notes: This table reports the niche-year regression for the Amazon ECS shock, separated from the three main AWS shocks reported in [Table 6](#) because the pre-trend test rejects. The treatment is the interaction of a post-shock indicator with a niche-level pre-period exposure score constructed from container and orchestration terminology in startup text. The entry outcome is log entrants and the exit outcome is the niche-year exit rate within six years. The specification includes niche fixed effects, year fixed effects, and controls for log media mentions. Standard errors clustered by niche are reported in parentheses. The pre-trend row reports the  $p$ -value from a clustered joint Wald test that all pre-period event-study coefficients equal zero. Statistical significance is \*(10%), \*\*(5%), or \*\*\*(1%).

**Figure D3:** Marketing keyword budget allocation survey

```
1 q2 = QuestionList(  
2   question_text = ""  
3   You were previously asked to read a company description and then draft  
4     10 non-branded Google Ads keyword search queries that would most  
5     effectively reach customers searching for a competing product or  
6     service.  
7  
8   Now you are being asked to make a budget allocation across the 10  
9   search queries (listed below). For each search query, select a  
10  float between 0 and 1 with EXACTLY two decimals.  
11  
12  The floats must sum to 1.00. Ties are allowed.  
13  If rounding to two decimals causes the numbers to not sum exactly to  
14  1.00 then assign the remainder to the FIRST  
15  keyword to make the total 1.00.  
16  Keep the floats in the SAME ORDER as the search queries.  
17  
18  Here are the search queries (in order): {{ marketing_keywords.answer  
19  }}  
20  
21  Here is the company description (for reference): "{{ scenario.  
22  company_description }}"  
23  """,  
24  min_list_items = 10,  
25  max_list_items = 10,  
26  question_name = "keywords_bid"  
27 )
```

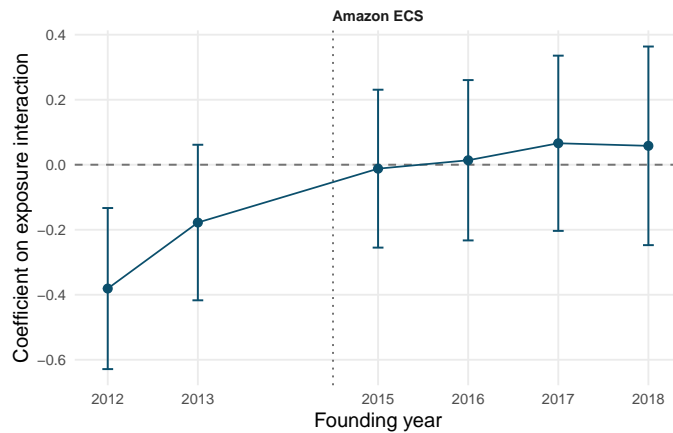
Notes: This figure shows the prompt that produces normalized weights for the marketing keywords the company would bid on to capture customer purchase intent. The portion “{{marketing\_keywords.answer}}” pipes the 10 keywords generated from the prior question described in [Figure D1a](#). We also update the model context with the company description by piping it in “{scenario.company\_description}”.

**Figure D4:** Concavity of Idea-Level Success in Entry



Notes: Each point represents the mean fraction of ideas with at least one exit within 6 years, averaged within 10 equal-frequency bins of total entry. The dashed curve plots the theoretical prediction  $P(n) = 1 - (1 - \hat{p})^n$  where  $\hat{p}$  is the sample mean firm-level exit rate. Under the model’s assumption of independent idiosyncratic failure, this function is concave in  $n$ .

**Figure D5:** Event-Study Evidence for the Amazon ECS Shock



Notes: This figure plots event-study coefficients for the Amazon ECS shock, normalizing the year before the post period to zero. Treatment intensity is the niche-level pre-period exposure score. The dashed horizontal line marks zero and the dotted vertical line marks the start of the post period. The visible pre-period movement and the joint Wald test ( $p = 0.011$ ) motivate separating this shock from the main results in [Figure 3](#).