

Learning What to Produce: Firm Dynamics over the Life Cycle*

Karam Jo[†]
Penn State

Seula Kim[‡]
Penn State

September 28, 2026

[\[Click Here for the Latest Version\]](#)

Abstract

Which young firms grow large and old? We show that the answer depends on how quickly firms learn what they are good at—their product-specific expertise. In our general-equilibrium model of firm dynamics, firms learn faster as they age and with better management and more data. Consistent with the theory, we find in U.S. Census data on manufacturers that young firms select products on estimated expertise about half as sharply as mature firms, and better-managed, more data-intensive firms select more sharply. In the calibrated model, firms that become large and old are mostly those born with high productivity, and among them, those that learn quickly what they are good at. Faster learning amplifies up-or-out dynamics: it raises early exit among firms with weak products but makes survivors more likely to become large. A policy that helps startups learn faster lowers survival but raises consumption by 0.5 percent net of cost, through better production decisions and selection inside firms.

JEL Code: D21, D22, D83, E23, M13, L11

Keywords: young firm dynamics, firm growth, Bayesian learning, product-specific expertise, information processing, product switching

*Previously circulated under the title “Product Switching and Young Firm Dynamics.” We are grateful to David Argente, Gadi Barlevy, Isaac Baley, Salome Baslandze (discussant), Andres Blanco, Nick Bloom, Johannes Boehm, Paco Buera, Satyajit Chatterjee, Gian Luca Clementi, Joel David, Maarten De Ridder (discussant), Jason Faberman, Alessandro Ferrari, Stefania Garetto, John Haltiwanger, Bart Hobijn, Hugo Hopenhayn, Boyan Jovanovic, Adrien Matray, Ryan Michaels, Javier Miranda, Alex Monge-Naranjo, Sara Moreira (discussant), Jane Olmstead-Rumsey, Marta Prato, Stephen Redding, Richard Rogerson, Leena Rudanko, Felipe Saffie, John Shea, Shihan Shen, Jim Tybout, Venky Venkateswaran, Steve Yeaple, and participants at various seminars and conferences, including the Atlanta Fed, Bank of Korea, Boston University TPRI, Brookings Institution, Chicago Fed, Kiel Institute, NYU, Philly Fed, Rice University, Sungkyunkwan University, University of Pittsburgh, U.S. Census Bureau, Washington University in St. Louis, ASSA, BSE Summer Forum, CEPR IMO-ESF Workshop, Econometric Society Meetings, Midwest Macro, SED, and SITE, for helpful comments. Any views expressed are those of the authors and not those of the U.S. Census Bureau. The Census Bureau has reviewed this data product to ensure appropriate access, use, and disclosure avoidance protection of the confidential source data used to produce this product. This research was performed at a Federal Statistical Research Data Center under FSRDC Project Number 2095. (CBDRB-FY25-P2095-R12268, CBDRB-FY26-P2095-R13144) All errors are our own.

[†]Email: karamjo@psu.edu. Address: 510 Kern Building, University Park, PA 16802.

[‡]Email: seulakim@psu.edu. Address: 614 Kern Building, University Park, PA 16802.

1 Introduction

Which young firms grow large and old? Most startups exit or stay small. Only a few grow quickly, but they account for much of the growth in aggregate employment and output.¹ In the canonical model of [Jovanovic \(1982\)](#), firms learn how good they are, and this learning determines which of them grow and which exit. We argue that firms also need to learn what they are good at, and that the answer depends on how quickly they do so.

A firm is born good at making and selling some products and poor at others. It learns which products match its strengths mainly through experience and drops those that turn out to be weak. Selection therefore starts inside the firm: just as workers discover their comparative advantage by working in different occupations, firms discover theirs by switching products.² We develop a theory of firm dynamics over the life cycle in which firms learn about their product-specific expertise—how well they can make and sell each product—and we test its core mechanism using product-level Census data on U.S. manufacturing firms.

The central idea is that firms are born with fundamental product-specific expertise, but this expertise is initially unknown. Firms differ in how precisely, and hence how quickly, they learn about it, with learning precision improving with age, managerial ability, and data. Young firms, and firms with weaker management or less data, form noisier beliefs: they tend to keep weak products too long and sometimes drop strong ones. Better-informed firms drop weak products sooner and continue searching for better ones, so learning amplifies the up-or-out dynamics of young firms ([Haltiwanger et al., 2013](#)).

In our calibrated model, the young firms that grow large and old are mostly those born with high productivity (as in [Sterk et al., 2021](#)), and among them, those that learn quickly what to produce. Faster learning cuts both ways: it raises early exit among firms with weak products but makes survivors grow and become larger with a better product mix.

Product switching is how firms act on what they learn, and it is common among U.S. manufacturers.³ In our data, about two-thirds of continuing firms add or drop a product

¹See [Dunne et al. \(1989\)](#), [Haltiwanger et al. \(2013\)](#), [Decker et al. \(2014\)](#), and [Haltiwanger et al. \(2017\)](#).

²On how workers learn about their ability, comparative advantage, and job-match quality, see [Jovanovic \(1979\)](#), [Papageorgiou \(2014\)](#), [Groes et al. \(2015\)](#), and [Baley et al. \(2022\)](#). [Hausmann and Rodrik \(2003\)](#) make a related point at the country level: countries develop by discovering what they are good at producing.

³See [Dunne et al. \(1988\)](#) and [Bernard et al. \(2010\)](#).

over five years, young and mature alike. However, switching matters more for young firms: it accounts for 56 percent of their growth, compared with 44 percent for mature firms.⁴ How well young firms choose their products may therefore shape which become large.

We make three contributions. The first is theoretical. In canonical models of firm learning, firms learn about a firm-level characteristic (Jovanovic, 1982; Arkolakis et al., 2018), while in models of product switching, firms either observe product-specific profitability (Bernard et al., 2010) or learn only about demand common across products (Timoshenko, 2015). We develop a general-equilibrium model with free entry in which each firm is born with three traits: a productive potential, which sets its productivity and maximum number of products; managerial ability; and permanent but unobserved expertise in each product.

Firms learn about this expertise from realized demand and an additional signal whose precision—their information-processing ability—rises with managerial ability, age, and accumulated data (Farboodi and Veldkamp, 2026; Baley and Veldkamp, 2025). Each period, firms choose how much to produce, which products to keep, and whether to search for new ones; they exit when no product is worth keeping. Because expertise is fixed at birth but revealed only over time, the model connects two views of which firms grow: one emphasizing heterogeneity present at birth (Sterk et al., 2021) and the other emphasizing learning after entry (Jovanovic, 1982).

The theory delivers three results. First, less precise information pulls beliefs toward the prior. Relative to known expertise, firms therefore underproduce products with sufficiently high expertise and overproduce those with sufficiently low expertise. Second, firms keep uncertain products partly for the option value of learning. Holding future precision fixed, a firm that knows more about a product gains less from this option, so it drops weak products sooner and searches more readily for new ones. Under cost conditions we characterize, higher productivity also encourages search but delays dropping by making weak products profitable enough to keep. Third, firm age matters beyond product tenure: a mature firm adding a new product starts from the same prior as a startup but learns faster.

⁴We measure growth as the five-year change in the shipment-weighted average of log real product shipments. Young firms also replace products more often: 39 percent both add and drop products within the same five years, compared with 33 percent of mature firms. Firms that add or drop products account for 91 percent of the output of continuing firms, consistent with Bernard et al. (2010).

Our second contribution is empirical: we test the model’s core prediction that firms keep the products they are good at, and do so more sharply when they process information more precisely. Using the Census of Manufactures from 1977 to 2012, we measure product-specific expertise as the persistent firm–product component of real shipments, as implied by the model’s demand structure, controlling for product–year conditions, firm–product tenure, and time-varying firm characteristics. Conditional on current shipments, a one-standard-deviation increase in estimated expertise is associated with a 26-percentage-point lower probability of dropping a product within five years among mature firms (older than five years), but only 14 points among young firms (aged five or less): selection on expertise is about half as strong in young firms. This gap remains when we use continuous firm age and when we estimate expertise from the same number of observations for each firm–product pair, so it does not reflect noisier estimates for young firms. Using the Management and Organizational Practices Survey (MOPS), we further show that selection on expertise is stronger among firms with more structured management, greater data availability, and more intensive use of data in decision making.

Our third contribution is quantitative. We calibrate the model to U.S. manufacturing, using the product-dropping regressions to discipline information-processing ability. Among entrants with high productivity, the probability of being active and large at age 20 ranges from 4 to 6 percent, depending on managerial ability; among all other entrants, it is below 0.3 percent.⁵ Among high-productivity entrants, those with high managerial ability are 4.6 percentage points less likely to survive to age 20 as they learn sooner that their products are weak. Conditional on survival, however, they are 5.9 points more likely to be large, so they are more likely to be both active and large (5.6 versus 4.4 percent). Like survival, current output is an incomplete guide to which young firms will succeed: among firms with similar output at age five, those with high productivity are 10 percentage points more likely to be active and in the top output decile ten years later.

Finally, we study a policy that improves how well new firms learn by raising their managerial capability at entry, without directly changing their productivity or expertise, at a

⁵Large firms have employment at or above the age-specific 90th percentile among active firms. Probabilities are measured relative to the original entrants.

resource cost of 0.1 percent of consumption. Following the same entrants with and without the subsidy, survival to age 20 falls from 20.1 to 18.0 percent, mainly because startups with weak initial products exit earlier. Yet cumulative real sales per entrant, counting zero sales after exit, rise by 2.8 percent, as better production and product choices more than offset the loss from lower survival. In general equilibrium, consumption rises by 0.5 percent net of the subsidy cost, even though the number of firms falls by 8 percent: better learning by startups leads to fewer but better firms.⁶ Evaluating policies for young firms by survival alone would therefore be misleading. In the model, the rationale for such programs is a knowledge spillover: startups do not internalize that their managerial knowledge benefits later cohorts (Romer, 1990; Jones, 1995).

Related Literature

First, we contribute to the literature on the sources of firm heterogeneity, selection, and young-firm growth. Young plants fail more often than older plants but grow faster if they survive (Dunne et al., 1989), and young firms' rapid productivity growth comes largely from selection and reallocation (Alon et al., 2018). Work on the sources of firm growth emphasizes differences present at birth, including demand for firms' products (Sedláček and Sterk, 2017), and founders' ambitions (Hurst and Pugsley, 2011; Guzman and Stern, 2020) and skills (Queiró, 2022). Sterk et al. (2021) find that such ex-ante heterogeneity accounts for much of the variation in firm size. Life-cycle growth also varies widely across countries (Hsieh and Klenow, 2014), in part because of weak selection and limits to delegation (Akcigit et al., 2021). Other work emphasizes frictions after entry, such as financing constraints (Cooley and Quadrini, 2001; Clementi and Hopenhayn, 2006; Schmalz et al., 2017), the gradual accumulation of customers and demand (Gourio and Rudanko, 2014; Foster et al., 2016; Argente et al., 2025a), technology diffusion (Akcigit and Ates, 2023; Jo and Kim, 2026), and workers' uncertainty about young employers (Kim, 2026). We show how a characteristic fixed at birth—a firm's expertise in each product—shapes growth only gradually as firms discover it through product switching, and how the speed of this discovery

⁶The phrase is from Ates and Saffie (2021), who show that tighter credit during sudden stops leads to fewer but more productive entrants. Here, selection instead reflects better information after entry.

determines both which firms grow and which exit.

Second, we contribute to studies on learning and firm dynamics. Canonical models study selection on a firm-level state (Jovanovic, 1982; Hopenhayn, 1992; Luttmer, 2007), while Arkolakis et al. (2018) link demand learning to firm growth, consistent with the exporter evidence in Berman et al. (2019). Empirical work also highlights the role of demand in selection and life-cycle growth (Foster et al., 2008; Eslava et al., 2024) and shows that firm size largely reflects product appeal and scope (Hottman et al., 2016). In Taiwanese textiles, differences in output quality and input efficiency across plants and products are the main source of productivity dispersion (Lee et al., 2026). We therefore move learning to the product level, so that firm-level selection emerges from product-level decisions, and we let the precision of learning differ across firms and over the life cycle.

Third, we contribute to the literature on product switching and firm growth. Models of innovation, reallocation, and distortions treat firms as collections of products (Klette and Kortum, 2004; Akcigit and Kerr, 2018; Acemoglu et al., 2018; Fattal Jaef, 2018). Product turnover is common and shapes firm growth (Bernard et al., 2010; Broda and Weinstein, 2010; Argente et al., 2018, 2024; Berlingieri et al., 2026); firms tend to add products that share inputs with their existing ones (Boehm et al., 2022); and standardizing products lets firms expand scope and scale at lower cost (Argente et al., 2025b).⁷ In Census data, existing plants also enter industries by changing their products (Dunne et al., 1988). Among entrants, such plants exit most often, and new plants of experienced firms least often (Dunne et al., 2005). The closest antecedents are Mitchell (2000), in which information links a firm's tasks and limits its scope, and Timoshenko (2015), in which new exporters switch products more often because they are still learning about demand common across products. We instead study learning about expertise in each product, show how it differs between young and mature firms, and trace its consequences for which firms grow large.

Finally, we relate to work on how information, management, and data shape firms' decisions. Imperfect information affects how firms allocate inputs and lowers aggregate productivity (David et al., 2016). Management practices are associated with productivity (Bloom

⁷In trade models, competition shifts production toward the products firms know they are best at (Eckel and Neary, 2010; Bernard et al., 2011; Mayer et al., 2014).

et al., 2019) and more accurate forecasts (Bloom et al., 2025), and field experiments show that improving management can raise firm performance (Bloom et al., 2013; Bruhn et al., 2018). Data can also improve the decisions of firms and their investors (Begenau et al., 2018; Farboodi et al., 2019; Farboodi and Veldkamp, 2026; Baley and Veldkamp, 2025). In Asriyan and Kohlhas (2026), firms acquire information to choose their scale, product type, and prices, and improvements in data processing raise TFP much more than welfare. We instead study how firms learn their expertise in each product over the life cycle, and how the speed of this learning shapes product retention, search, and exit.

Related evidence shows that startups that test their ideas more systematically perform better and pivot more often (Camuffo et al., 2020), and are more likely both to scale and to fail quickly (Koning et al., 2022), in line with the view of entrepreneurship as experimentation (Kerr et al., 2014). Our model generates the same pattern for firms that learn faster. As Kehrig and Vincent (2026) emphasize in a different setting, dispersion generated by noisy learning need not reflect inefficiency. Our policy result also relates to Acemoglu et al. (2018), who find that subsidies to incumbent R&D help weak firms survive, whereas taxing continued operation by incumbents raises welfare.⁸

The rest of the paper proceeds as follows. Section 2 presents the model and its theoretical results. Section 3 describes the data, constructs our measure of expertise, and tests the model’s core mechanism. Section 4 calibrates the model and studies selection within and across firms over the life cycle, a subsidy to startups’ managerial investment, and welfare. Section 5 concludes.

2 Theoretical Framework

This section presents the model. Firms learn about their expertise in each product, and their information-processing ability depends on managerial ability, firm age, and accumulated data, so firms learn at different rates. The model builds on learning about firm fundamentals (Jovanovic, 1982; Kim, 2026) and on product switching (Bernard et al., 2010), and the

⁸Our welfare analysis also relates to work on knowledge accumulation and diffusion (Romer, 1990; Jones, 1995; Perla and Tonetti, 2014; Sampson, 2016; Buera and Oberfield, 2020) and on the transfer of managerial and entrepreneurial knowledge (Giorcelli, 2019; Guiso et al., 2021).

role of information processing is motivated by work on the data economy (Farboodi et al., 2019; Farboodi and Veldkamp, 2026; Baley and Veldkamp, 2025). We first describe the environment, consumers, and firms, then potential startups and the equilibrium, and finally derive the model's implications.

2.1 Environment

The economy contains a finite set of products $p \in \{1, \dots, N\}$ and a continuum of firms indexed by i . Product markets are monopolistically competitive, with firms producing differentiated goods within each market. Time is discrete, indexed by t .

2.2 Consumers

A representative household supplies its labor endowment $\bar{L}$ inelastically and owns all firms. It maximizes

$$\sum_{t=0}^{\infty} \beta^t \log \bar{C}_t,$$

subject to the period budget constraint $\sum_{p=1}^N \int_{i \in \Omega_{pt}} p_{ipt} c_{ipt} di + w_t \mathcal{F}_t^e \leq w_t \bar{L} + \Pi_t$, where $\beta \in (0, 1)$ is the discount factor. Here, Ω_{pt} is the set of firms operating in product market p , p_{ipt} and c_{ipt} are the price and consumption of firm i 's good in market p , w_t is the wage, and Π_t denotes aggregate profits before entry costs. Aggregate entry costs, including pre-entry capability investments, are $\mathcal{F}_t^e$ in labor units.

The aggregate consumption index, $\bar{C}_t$, and product-level consumption indices, C_{pt} , are defined by nested constant elasticity of substitution (CES) aggregators:

$$\bar{C}_t = \left(\sum_{p=1}^N C_{pt}^{\frac{\nu-1}{\nu}} \right)^{\frac{\nu}{\nu-1}}, \quad C_{pt} = \left[\int_{i \in \Omega_{pt}} (\theta_{ipt} c_{ipt})^{\frac{\eta-1}{\eta}} di \right]^{\frac{\eta}{\eta-1}}.$$

The elasticities of substitution across products and differentiated goods within each market are ν and η , respectively, with $1 < \nu < \eta$. The firm-product demand shifter θ_{ipt} combines permanent product-specific expertise, θ_{ip} , with a transitory demand shock, as specified be-

low. Expenditure minimization yields demand for each good:⁹

$$c_{ipt} = P_{pt}^\eta C_{pt} \theta_{ipt}^{\eta-1} p_{ipt}^{-\eta}. \quad (2.1)$$

The corresponding aggregate and product-level price indices are

$$\bar{P}_t = \left(\sum_{p=1}^N P_{pt}^{1-\nu} \right)^{\frac{1}{1-\nu}}, \quad P_{pt} = \left[\int_{i \in \Omega_{pt}} \theta_{ipt}^{\eta-1} p_{ipt}^{1-\eta} di \right]^{\frac{1}{1-\eta}}.$$

2.3 Firms

Before entry, potential startups invest in productivity and managerial capabilities, summarized by z_0 and m_0 , respectively. Higher z_0 raises productivity, while higher m_0 improves information-processing ability and accelerates learning about product-specific expertise. We return to these pre-entry choices in Section 2.4.

Firms are born with two permanent fundamental types: productive potential φ_i and managerial ability m_i . Productive potential determines productivity $z_i \equiv z(\varphi_i; z_0)$ and potential product scope $n_i \equiv n(\varphi_i)$, the maximum number of product lines a firm can operate simultaneously. Productivity is strictly increasing in φ_i , while potential scope is nondecreasing. Managerial ability m_i may be correlated with φ_i and is drawn from a distribution shifted by m_0 . In each period, firms decide for each product line whether to add or drop a product and choose quantities for all operating products.

2.3.1 Firm-Product Expertise and Learning

In addition to its fundamental types (φ_i, m_i) , firm i is born with time-invariant expertise θ_{ip} in each product p , drawn from a lognormal distribution, $\log \theta_{ip} \sim N(\theta_0, \sigma_0^2)$. Expertise captures how well the firm can meet consumer needs in a particular product market, including preferences for quality, design, and functionality. We model expertise as a demand-side attribute that raises consumers' valuation of the firm's good and its profitability.¹⁰

⁹The inverse demand function is $p_{ipt} = P_{pt} C_{pt}^{1/\eta} \theta_{ipt}^{(\eta-1)/\eta} c_{ipt}^{-1/\eta}$.

¹⁰Alternatively, expertise can be modeled on the production side as product-specific productivity.

Conditional on firm types, expertise draws are independent across products. In the baseline, market and additional-signal errors are mutually independent across products and dates and independent of expertise.

Firms do not observe θ_{ip} directly. After production and sales, they observe a noisy signal inferred from the market-clearing price induced by realized demand:

$$\log \theta_{ipt} = \log \theta_{ip} + \varepsilon_{ipt},$$

where $\varepsilon_{ipt} \sim N(0, \sigma_\varepsilon^2)$ is an i.i.d. demand shock. Learning about θ_{ip} follows Bayesian updating based on $\log \theta_{ipt}$ and an additional signal $\log s_{ipt}$ obtained by combining realized market performance with other information, such as knowledge of customer tastes and market conditions. Firms with greater experience, data, and managerial ability can extract a more precise signal from this information (Farboodi et al., 2019; Farboodi and Veldkamp, 2026; Baley and Veldkamp, 2025; Bloom et al., 2025).¹¹ We model this additional signal as

$$\log s_{ipt} = \log \theta_{ip} + v_{ipt}, \quad v_{ipt} \sim N\left(0, \frac{\sigma_\varepsilon^2}{\psi_{it}}\right). \quad (2.2)$$

We define ψ_{it} as the firm's information-processing precision, with higher ψ_{it} implying a more precise additional signal. A firm's ability to process observable information depends on its managerial ability m_i , experience (captured by firm age a_{it}), and accumulated data D_{it} , which depend on a_{it} , potential product scope n_i , and potentially other data sources. We summarize these factors by $\psi_{it} \equiv \psi(m_i, a_{it}, D_{it}; \varphi_i)$, with $\partial\psi_{it}/\partial x \geq 0$ for $x \in \{m_i, a_{it}, D_{it}\}$.¹² Conditional on (φ_i, m_i) , a_{it} determines the evolution of ψ_{it} .

Given this structure, firm i enters market p with prior $\log \theta_{ip} \sim N(\theta_0, \sigma_0^2)$. After production and sales, it observes the market outcome and additional signal and updates its beliefs about θ_{ip} . After $\tau_{ipt} \geq 1$ production periods, including t , its information set is

¹¹An alternative interpretation is that firms use their information about customer tastes and market conditions to filter transitory demand fluctuations from realized market performance.

¹²The cumulative number of products dropped can serve as another proxy for firm experience. We abstract from this variable because it introduces dependencies across product lines, reducing tractability.

$\mathcal{I}_{ipt} \equiv \{\theta_{ipt}, s_{ipt}\}_{\tilde{t}=t-\tau_{ipt}+1}^t$. This yields the posterior mean and variance of $\log \theta_{ip}$:

$$\hat{\theta}_{ipt} = \hat{\sigma}_{ipt}^2 \left[\frac{\theta_0}{\sigma_0^2} + \frac{\tau_{ipt}}{\sigma_\varepsilon^2} \left(\frac{1}{\tau_{ipt}} \sum_{\tilde{t}=t-\tau_{ipt}+1}^t \log \theta_{ipt\tilde{t}} \right) + \frac{\sum_{\tilde{t}=t-\tau_{ipt}+1}^t \psi_{it\tilde{t}} \log s_{ipt\tilde{t}}}{\sigma_\varepsilon^2} \right], \quad (2.3)$$

$$\hat{\sigma}_{ipt}^2 = \left[\frac{1}{\sigma_0^2} + \frac{\tau_{ipt}}{\sigma_\varepsilon^2} + \frac{\sum_{\tilde{t}=t-\tau_{ipt}+1}^t \psi_{it\tilde{t}}}{\sigma_\varepsilon^2} \right]^{-1} = \left[\frac{1}{\hat{\sigma}_{ip,t-1}^2} + \frac{1 + \psi_{it}}{\sigma_\varepsilon^2} \right]^{-1}. \quad (2.4)$$

The posterior mean is a weighted average of the prior mean θ_0 , the average observed outcome, and additional signals, with weights determined by firm-product tenure τ_{ipt} and the history of information-processing ability $\psi_{it\tilde{t}}$. A central feature of the model is the distinction between firm-product tenure and information-processing ability; Lemma 1 shows how the latter shapes posterior beliefs at a given tenure. The posterior also determines the predictive distribution of $\theta_{ip,t+1}$.

Since posterior variance is determined by firm-product tenure, firm age, and fundamental types, posterior beliefs can be tracked by $(\hat{\theta}_{ipt}, \tau_{ipt}; a_{it}, \psi_{it}, \varphi_i, m_i)$. We retain ψ_{it} explicitly to emphasize the role of information-processing ability. All firms start with the common prior $(\hat{\theta}_{ipt}^0, (\hat{\sigma}_{ipt}^0)^2) \equiv (\theta_0, \sigma_0^2)$ at tenure $\tau_{ipt}^0 = 0$.

Lemma 1 (Biased Learning). *At fixed product tenure $\tau \geq 1$, lower information-processing ability ψ_{it} , holding earlier processing fixed, increases posterior variance and magnifies expected underestimation of $\log \theta_{ip}$ above θ_0 and expected overestimation below θ_0 .¹³*

Proof. See Appendix A.1. □

These conditional biases reflect Bayesian shrinkage: lower cumulative processing over the observation window gives the prior greater weight.

Firm-product tenure and firm age affect learning through distinct channels. Longer tenure τ provides more product-specific observations, whereas greater information-processing ability ψ_{it} yields more precise additional signals. Conditional on (φ_i, m_i) , ψ_{it} increases with firm age. Thus, when young and mature firms start a product with the same prior, the mature firm can learn faster, with lower posterior variance and weaker shrinkage

¹³Expectations average over signal histories given true expertise, without conditioning on retention.

toward the prior at the same product tenure. Differences in managerial ability m_i generate learning heterogeneity among firms of the same age and product tenure. These differences lead to systematically different product-switching decisions over the life cycle.

Lastly, we characterize the transition of posterior beliefs conditional on continued production. Under our functional-form assumption, ψ_{it} is determined by (a_{it}, φ_i, m_i) , so both $\hat{\sigma}_{ipt}^2$ and $\psi_{i,t+1}$ are known given the current state. Continued production implies $\tau_{ip,t+1} = \tau_{ipt} + 1$ and $a_{i,t+1} = a_{it} + 1$. After observing the market outcome and additional signal in period $t + 1$, Bayesian updating gives $\hat{\sigma}_{ip,t+1}^2 = (1/\hat{\sigma}_{ipt}^2 + (1 + \psi_{i,t+1})/\sigma_\varepsilon^2)^{-1}$ and $\hat{\theta}_{ip,t+1} = \hat{\sigma}_{ip,t+1}^2 \left(\hat{\theta}_{ipt}/\hat{\sigma}_{ipt}^2 + \log \theta_{ip,t+1}/\sigma_\varepsilon^2 + \psi_{i,t+1} \log s_{ip,t+1}/\sigma_\varepsilon^2 \right)$.

Integrating over $\log \theta_{ip}$ using the current posterior and over the independent signal errors yields a Gaussian transition for the posterior mean.¹⁴ The transition kernel $g_{it}(\cdot \mid \hat{\theta}_{ipt}, \tau_{ipt}) \equiv g(\cdot \mid \hat{\theta}_{ipt}, \tau_{ipt}; a_{it}, \psi_{it}, \varphi_i, m_i)$ is

$$g_{it}(\hat{\theta}_{ip,t+1}, \tau_{ip,t+1} \mid \hat{\theta}_{ipt}, \tau_{ipt}) = \phi \left(\hat{\theta}_{ip,t+1}; \hat{\theta}_{ipt}, \hat{\sigma}_{ipt}^2 - \hat{\sigma}_{ip,t+1}^2 \right) \mathbf{1}\{\tau_{ip,t+1} = \tau_{ipt} + 1\}, \quad (2.5)$$

where ϕ denotes the normal density, parameterized by its mean and variance. The transition variance equals the reduction in posterior uncertainty, while the expected next posterior mean equals its current value. Hence, the posterior-mean process is a martingale under the firm's predictive beliefs.

2.3.2 Timing of the Model

Figure 1 illustrates the timing. Firms begin period t with posterior beliefs about θ_{ip} and predetermined quantities, both based on $\mathcal{I}_{ipt-1}$, prior to the realization of θ_{ipt} .¹⁵ Demand and market-clearing prices are then realized given θ_{ipt} , followed by exogenous death with probability δ . Remaining firms update their beliefs using realized outcomes and additional signals and decide which products to retain or drop. If all active products are dropped, the firm exits before search; otherwise, it can search for new products.

¹⁴The next posterior mean is conditionally Gaussian because it is an affine function of the jointly Gaussian $\log \theta_{ip}$ and signal errors. By the law of total variance, its conditional variance equals the current posterior variance minus the next posterior variance, which is known given the current state.

¹⁵Entrants begin with the prior $(\hat{\theta}_{ip}^0, (\hat{\sigma}_{ip}^0)^2)$ and observe m_i after initial production, before their first belief update.




Figure 1: Timeline of the model

Potential startups make entry decisions concurrently with incumbents' search decisions. Search outcomes and entrants' product assignments are realized at the end of the period. Continuing firms choose next-period quantities using updated beliefs for retained products, while beliefs for new products reset to the initial prior upon successful switching into a new product market. We define the value function after the death shock and belief updating, before product-retention and search decisions.

2.3.3 Production

Firms produce each product using linear technology with labor as the only input, $y_{ipt} = z_i \ell_{ipt}$, where $z_i = z(\varphi_i; z_0)$. Each operating product incurs a fixed cost f^o in labor units. Quantities are chosen using posterior beliefs about θ_{ip} before θ_{ipt} is realized. In this subsection, expectations are conditional on the current state $(\hat{\theta}_{ipt}, \tau_{ipt}; a_{it}, \psi_{it}, \varphi_i, m_i)$, and firms take $(P_{pt+1}, C_{pt+1}, w_{t+1})$ as given and known. For a product continuing into period $t+1$, the firm maximizes expected per-period profit:

$$\bar{\pi}_{it+1}(\hat{\theta}_{ipt}, \tau_{ipt}) \equiv \max_{y_{ipt+1} \geq 0} \left\{ \mathbb{E}_{\theta_{ipt+1}} [p_{ipt+1}(y_{ipt+1}, \theta_{ipt+1})] y_{ipt+1} - w_{t+1} \left(\frac{y_{ipt+1}}{z_i} + f^o \right) \right\}, \quad (2.6)$$

subject to demand (2.1). The optimal quantity is

$$y_{ipt+1} = \left(\frac{\eta - 1}{\eta} \right)^\eta P_{pt+1}^\eta C_{pt+1} \left[\mathcal{B}_{it}(\hat{\theta}_{ipt}, \tau_{ipt}) \right]^\eta \left(\frac{z_i}{w_{t+1}} \right)^\eta, \quad (2.7)$$

where $\mathcal{B}_{it}(\hat{\theta}_{ipt}, \tau_{ipt}) \equiv \mathcal{B}(\hat{\theta}_{ipt}, \tau_{ipt}; a_{it}, \psi_{it}, \varphi_i, m_i)$ is the conditional expectation of $\theta_{ipt+1}^{(\eta-1)/\eta}$:

$$\mathcal{B}_{it}(\hat{\theta}_{ipt}, \tau_{ipt}) = \exp \left\{ \frac{\eta-1}{\eta} \hat{\theta}_{ipt} + \frac{1}{2} \left(\frac{\eta-1}{\eta} \right)^2 (\hat{\sigma}_{ipt}^2 + \sigma_\varepsilon^2) \right\}. \quad (2.8)$$

Substituting (2.7) into (2.1) yields the market-clearing price after θ_{ipt+1} is realized:

$$p_{ipt+1} = \frac{\eta}{\eta-1} \frac{w_{t+1}}{z_i} \frac{\theta_{ipt+1}^{(\eta-1)/\eta}}{\mathcal{B}_{it}(\hat{\theta}_{ipt}, \tau_{ipt})}. \quad (2.9)$$

Thus, the expected price is a constant markup over marginal cost, $\mathbb{E}_{\theta_{ipt+1}}[p_{ipt+1}] = \frac{\eta}{\eta-1} \frac{w_{t+1}}{z_i}$.

Combining this with (2.7) gives optimal expected per-period profit:

$$\bar{\pi}_{it+1}(\hat{\theta}_{ipt}, \tau_{ipt}) = \frac{(\eta-1)^{\eta-1}}{\eta^\eta} \left(\frac{z_i}{w_{t+1}} \right)^{\eta-1} P_{pt+1}^\eta C_{pt+1} \left[\mathcal{B}_{it}(\hat{\theta}_{ipt}, \tau_{ipt}) \right]^\eta - w_{t+1} f^o. \quad (2.10)$$

The quantity and price equations motivate the empirical specification for estimating firm-product expertise in (3.40).¹⁶

2.3.4 Product Switching and Exit

Each firm i is endowed with n_i product lines, indexed by $k \in \{1, \dots, n_i\}$, representing capacity slots to which products can be assigned over time. Let $p_{kt} \in \{1, \dots, N, \emptyset\}$ denote the product assigned to line k in period t , with $p_{kt} = \emptyset$ for a vacant line. We write p_k when the time index is clear. Product assignments evolve endogenously through firms' product-switching decisions. When a firm searches, a destination product is randomly drawn for the line. With probability x , search succeeds and the destination replaces the existing product

¹⁶The model counterpart of real shipments is $p_{ipt} y_{ipt} / P_{pt}$. Combining (2.7) and (2.9) gives

$$\begin{aligned} \log \left(\frac{p_{ipt} y_{ipt}}{P_{pt}} \right) &= (\eta-1) \log \left(\frac{\eta-1}{\eta} \right) + (\eta-1) \log P_{pt} + \log C_{pt} \\ &\quad + (\eta-1) \log \left(\frac{z_i}{w_t} \right) + (\eta-1) \log \mathcal{B}_{ip,t-1} + \frac{\eta-1}{\eta} (\log \theta_{ip} + \varepsilon_{ipt}), \end{aligned} \quad (2.11)$$

where $\mathcal{B}_{ip,t-1} \equiv \mathcal{B}_{i,t-1}(\hat{\theta}_{ip,t-1}, \tau_{ip,t-1})$. Product-year effects absorb the aggregate terms, while firm and firm-product controls account for other sources of shipment variation. Appendix B.2 discusses the mapping from estimated fixed effects to structural expertise.




Figure 2: Evolution of product lines for firm i with potential product scope $n_i = n(\varphi_i)$

or fills a vacancy.¹⁷ Product lines are therefore the units of capacity and decision-making, while product indices identify markets.

A product line's value depends on the firm's posterior beliefs about product-specific expertise. Given operating costs and a fixed search cost f^s in labor units per searching line, the firm faces two extensive-margin decisions per line: whether to retain or drop its product and whether to search for another. The firm may retain a product while searching for its replacement. If all active products are dropped, the firm exits before search; otherwise, it searches to fill newly and previously vacant lines. Firms exit exogenously with probability δ . We characterize search and drop thresholds below.

Figure 2 illustrates these decisions and outcomes for firm i with $n_i = n(\varphi_i)$. Bar heights represent $\mathcal{B}_{it}(\hat{\theta}, \tau)$ after belief updating but before switching in panel (a), and after search outcomes in panel (b). Different patterns identify distinct products, while vacant lines have no bar. The dashed line indicates $\mathcal{B}_{it}(\hat{\theta}^0, 0)$ under the common prior.

Products in lines 1 and 2 lie below their drop thresholds and are dropped. The product in line 3 lies above its search threshold and is retained without search. Products in lines k and $n_i - 1$ lie between their respective drop and search thresholds and are retained while searches are conducted. Because active products remain, the firm also searches to fill newly

¹⁷The success probability x is analogous to R&D productivity in endogenous growth models. Firms choose search participation rather than intensity; the fixed cost f^s and x jointly determine participation and expected product additions.

vacant lines 1 and 2 and the initially vacant line n_i .

Searches succeed for lines 1, k , and n_i , which receive new products with beliefs reset to the common prior. Searches fail for lines 2 and $n_i - 1$: line 2 remains vacant, while line $n_i - 1$ retains its existing product. An initially vacant line would also remain vacant after failed search, a case not shown in the figure.

2.3.5 Value Function of Product Lines

Since product-switching decisions are made within product lines, we formulate the recursive problem at the product-line level. Endogenous exit, however, can link a line's continuation value to posterior beliefs about the firm's other products through the exit probability. We establish the conditions for portfolio independence in Section 2.3.7. We consider firms aged two or older; Section 2.4 presents the age-one startup problem.

Let $V_{it}^{(\cdot)}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) \equiv V^{(\cdot)}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}; a_{it}, \psi_{it}, \varphi_i, m_i)$. We likewise suppress firm-level states in beliefs, thresholds, and transition probabilities when there is no ambiguity. For an occupied product line k , the value function is

$$V_{it}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) = \max\{V_{it}^{ns}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}), V_{it}^{cs}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}), V_{it}^{ds}(\emptyset)\}, \quad (2.12)$$

where V_{it}^{ns} is the value of retaining the product without search, V_{it}^{cs} is the value of retaining it while searching, and $V_{it}^{ds}(\emptyset)$ is the value of dropping it and searching, conditional on firm continuation. If all active products are dropped, the firm exits before search; the associated exit payoff is specified below. The corresponding search and drop thresholds, $\mathcal{B}_{it}^{cs}(\tau_{ip_k t})$ and $\mathcal{B}_{it}^{ds}(\tau_{ip_k t})$, are characterized in Proposition 1.¹⁸

(a) No search ($\mathcal{B}_{it}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) \geq \mathcal{B}_{it}^{cs}(\tau_{ip_k t})$): The firm retains the product without searching this period but may search later as beliefs evolve. The corresponding value is

$$V_{it}^{ns}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) = \beta \left[\bar{\pi}_{it+1}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) + (1 - \delta) \mathbb{E}_{ip_k t} V_{it+1}(\hat{\theta}_{ip_k t+1}, \tau_{ip_k t+1}) + \delta V_{it+1}^\delta \right], \quad (2.13)$$

¹⁸We focus on continuing firms that search in vacant lines, with the search value defined in (2.15). Proposition 1 gives conditions for finite, strictly ordered thresholds.

where $\mathbb{E}_{ip_k t}$ conditions on the current product-line and firm-level states, and V_{it+1}^δ is the terminal value upon exogenous firm exit, which we normalize to zero. The expected continuation value incorporates next period's optimal choices in (2.12):

$$\mathbb{E}_{ip_k t} V_{it+1}(\hat{\theta}_{ip_k t+1}, \tau_{ip_k t+1}) = \int_{\mathbb{R}} V_{it+1}(\hat{\theta}', \tau') g_{it}(\hat{\theta}', \tau' \mid \hat{\theta}_{ip_k t}, \tau_{ip_k t}) d\hat{\theta}',$$

where $\tau' = \tau_{ip_k t} + 1$ and g_{it} is defined in (2.5). The corresponding decision regions are

$$\begin{aligned} \mathcal{R}_{ip_k t+1}^{ns} &= \{(\hat{\theta}', \tau') : \mathcal{B}_{it+1}(\hat{\theta}', \tau') \geq \mathcal{B}_{it+1}^{cs}(\tau')\}, \\ \mathcal{R}_{ip_k t+1}^{cs} &= \{(\hat{\theta}', \tau') : \mathcal{B}_{it+1}^{ds}(\tau') \leq \mathcal{B}_{it+1}(\hat{\theta}', \tau') < \mathcal{B}_{it+1}^{cs}(\tau')\}, \\ \mathcal{R}_{ip_k t+1}^{ds} &= \{(\hat{\theta}', \tau') : \mathcal{B}_{it+1}(\hat{\theta}', \tau') < \mathcal{B}_{it+1}^{ds}(\tau')\}. \end{aligned}$$

The last region corresponds to dropping the product and includes endogenous exit when all active products are dropped.

(b) Continue/Switch with search ($\mathcal{B}_{it}^{ds}(\tau_{ip_k t}) \leq \mathcal{B}_{it}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) < \mathcal{B}_{it}^{cs}(\tau_{ip_k t})$): In this region, the firm searches for a new product while retaining the current product. If search succeeds, the new product replaces it at the end of the period; otherwise, the firm retains it. Using the no-search value in (2.13), we write

$$V_{it}^{cs}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) = -w_t f^s + x V_{it}^{ns}(\hat{\theta}_{ip'_k t}^0, 0) + (1-x) V_{it}^{ns}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}), \quad (2.14)$$

where p'_k denotes the new destination product. Upon success, we treat the firm as having no prior experience in the destination market and reset its beliefs to the common prior, with $\hat{\theta}_{ip'_k t}^0 = \theta_0$ and tenure zero.¹⁹ Following failed search, the payoff coincides with the no-search value because the firm retains the existing product.

Search operates as follows. The firm's exclusion set contains its current products and the most recent destination of each vacant line, yielding $n(\varphi_i)$ distinct excluded products.²⁰

¹⁹For tractability, we apply the prior-reset convention whenever a destination product is adopted. If the firm returns to a previously operated product, earlier information could be retained; we abstract from this history dependence. Our calibration uses a large product space, $N = 1400$.

²⁰We assume $n(\varphi_i) \leq \lfloor N/2 \rfloor$, ensuring enough distinct destinations for all product lines.

Destinations are drawn uniformly without replacement across searching lines from the remaining $N - n(\varphi_i)$ products. Search succeeds independently across lines with probability x , in which case the destination is adopted. If search on a vacant line fails, the unsuccessful draw becomes its most recent destination for the next search. Under uniform destination draws and symmetric product markets, expected continuation values are identical across feasible destinations, allowing us to suppress expectations over destination products.²¹

(c) Drop/Switch with search ($\mathcal{B}_{it}(\hat{\theta}_{ip_k t}, \tau_{ip_k t}) < \mathcal{B}_{it}^{ds}(\tau_{ip_k t})$): In this case, the firm drops the current product and, conditional on firm continuation, searches to fill the vacant line. If search succeeds, the line is filled with the new product; otherwise, it remains vacant. Using (2.13), the value is

$$V_{it}^{ds}(\emptyset) = -w_t f^s + x V_{it}^{ns}(\hat{\theta}_{ip_k t}^0, 0) + (1 - x)\beta [(1 - \delta)V_{it+1}^{ds}(\emptyset) + \delta V_{it+1}^\delta]. \quad (2.15)$$

After failed search, the focal line remains vacant. We assume that its payoff upon endogenous firm exit equals its value under firm continuation, $V_{it+1}^{ds}(\emptyset)$. Thus, the endogenous exit probability drops out of the continuation value. Section 2.3.7 formalizes this payoff-equivalence assumption.

2.3.6 Search and Product-Drop Thresholds

We characterize the product-line decisions of firms aged two or older, taking stationary aggregate prices and quantities as given. We suppress the product index and common firm-level states when there is no ambiguity.

At the search threshold, retaining the product with and without search yields the same value. Equation (2.14) gives

$$V_{it}^{ns}(\hat{\theta}^0, 0) - V_{it}^{ns}(\hat{\theta}_{it}^{cs}(\tau), \tau) = \frac{w_t f^s}{x}. \quad (2.16)$$

The value gain from a successful replacement therefore equals the search cost per successful

²¹Appendix A.3 shows that, despite sampling without replacement across searching lines, each line is marginally assigned any feasible destination with probability $1/[N - n(\varphi_i)]$.

search. At the drop threshold, retaining the product while searching yields the same value as dropping it:

$$V_{it}^{cs}(\hat{\theta}_{it}^{ds}(\tau), \tau) = V_{it}^{ds}(\emptyset). \quad (2.17)$$

The common search cost and successful-search payoff cancel in this comparison; the firm compares retaining the incumbent with a vacancy following failed search.

Proposition 1 (Existence, Uniqueness, and Ordering of Thresholds). *Fix a firm type and age. Suppose $f^o > 0$, $0 < x < 1$, and the vacancy values satisfy*

$$V_{is}^{ds}(\emptyset) > \beta(1 - \delta)V_{i,s+1}^{ds}(\emptyset), \quad s = t, t + 1, \dots \quad (2.18)$$

For each feasible product tenure $\tau \geq 1$, there exist unique finite cutoffs $\hat{\theta}_{it}^{ds}(\tau) < \hat{\theta}_{it}^{cs}(\tau)$ satisfying (2.17) and (2.16). The optimal policy is to drop the product below $\hat{\theta}_{it}^{ds}(\tau)$, retain it while searching between the cutoffs, and retain it without search above $\hat{\theta}_{it}^{cs}(\tau)$. Since $\mathcal{B}_{it}(\hat{\theta}, \tau)$ is strictly increasing in $\hat{\theta}$, these cutoffs uniquely define $0 < \mathcal{B}_{it}^{ds}(\tau) < \mathcal{B}_{it}^{cs}(\tau) < \infty$.

Proof. See Appendix A.2. □

Condition (2.18) requires a vacant line's value today to exceed its discounted value next period. Both thresholds are crossings of the same increasing no-search value with different levels; this condition orders those levels and ensures a nonempty continue-and-search region. Online Appendix OA.A.3 gives a sufficient condition in terms of prior expected operating profit, search cost, and search success probability. Adopting a new product resets product-specific beliefs and tenure while preserving firm age and the information-processing path. Section 2.4 treats age-one startups' retention and exit decisions separately.

2.3.7 Firm Value Function

Let $\Theta_{it} \equiv \{(\hat{\theta}_{ipkt}, \tau_{ipkt})\}_{k=1}^{n_i}$ denote firm i 's portfolio of product-line states, with $\emptyset$ denoting a vacant line. If all active products are dropped, the firm exits before search. For a firm operating a single product, dropping that product is therefore an exit decision. Search on dropped or vacant lines therefore depends on retention elsewhere in the firm. For firms aged two or older, we assume that each line receives $V_{it}^{ds}(\emptyset)$ upon endogenous exit, equal

to its vacant-line value under firm continuation.²² Equating these payoffs removes that dependence from line values, yielding product-line separability.

Lemma 2 (Product-Line Separability). *For $a_{it} \geq 2$, under payoff equivalence and the maintained search and learning assumptions, a product line's value depends only on its own state $(\hat{\theta}_{ip_k t}, \tau_{ip_k t})$ and the common firm-level states $(a_{it}, \psi_{it}, \varphi_i, m_i)$.*

Proof. See Appendix A.3. □

Operating profits and operating and search costs are additive across product lines, with no firm-level fixed operating cost.

Proposition 2 (Additive Firm Value). *For $a_{it} \geq 2$, under the assumptions of Lemma 2,*

$$V(\Theta_{it}; a_{it}, \psi_{it}, \varphi_i, m_i) = \sum_{k=1}^{n_i} V(\hat{\theta}_{ip_k t}, \tau_{ip_k t}; a_{it}, \psi_{it}, \varphi_i, m_i),$$

where a vacant line contributes $V^{ds}(\emptyset; a_{it}, \psi_{it}, \varphi_i, m_i)$.

Proof. See Appendix A.3. □

2.4 Potential Startups

Each period, potential startups make entry and capability-investment decisions concurrently with incumbents' search decisions. Pre-entry investments (a_t^m, a_t^z) determine managerial and productivity capabilities, m_{0t} and z_{0t} , according to

$$\log m_{0t} = (1 - \zeta_m) \log a_t^m + \zeta_m \log M_{t-1}, \quad \log z_{0t} = (1 - \zeta_z) \log a_t^z + \zeta_z \log Z_{t-1}. \quad (2.19)$$

Here, M_{t-1} and Z_{t-1} denote aggregate managerial and productivity capabilities inherited from the previous cohort, with $M_{t-1} = m_{0,t-1}$ and $Z_{t-1} = z_{0,t-1}$ in equilibrium. $\zeta_m, \zeta_z \in [0, 1)$ govern inter-cohort transferability.²³

²²The payoff upon exogenous exit remains zero, $V^\delta = 0$.

²³In the spirit of “standing on the shoulders of giants,” this specification combines current investment with inherited knowledge (Romer, 1990; Jones, 1995). Evidence on entrepreneurial and managerial knowledge transmission motivates $\zeta_m > 0$ (Guiso et al., 2021; Bloom et al., 2019; Giorcelli, 2019), while models of idea diffusion motivate $\zeta_z > 0$ (Perla and Tonetti, 2014; Sampson, 2016; Buera and Oberfield, 2020).

The startup resource cost, in labor units, is $f_t^e(a_t^m, a_t^z) = \bar{f}^e + \alpha_m^e (a_t^m)^\chi + \alpha_z^e (a_t^z)^\chi$, where $\bar{f}^e$ is the fixed startup cost, α_m^e and α_z^e scale investment costs, and χ governs their curvature. Potential startups take these parameters and inherited capabilities as given. For the entering cohort, m_{0t} shifts the distribution of managerial ability m_i , while z_{0t} determines productivity through $z_i = z_{0t}\varphi_i$.

Upon entry, startups draw permanent fundamental types (φ_i, m_i) , with $\varphi_i \sim H_\varphi$ and $m_i \mid \varphi_i \sim H_m(\cdot \mid \varphi_i; m_{0t})$. The firm observes its productive potential φ_i during production planning and decides whether to operate.²⁴ Entry is undirected: each startup receives one initial product market, with its remaining $n(\varphi_i) - 1$ product lines vacant.²⁵ Conditional on operating, the firm produces and sells its initial product in $t + 1$. Managerial ability m_i , unknown during entry and initial production, is revealed before the first Bayesian update about product-specific expertise.

After the first update, the age-one startup chooses whether to retain its initial product and whether to search for its replacement. Retention also allows the firm to search in its $n(\varphi_i) - 1$ vacant lines; dropping the initial product triggers exit before search.²⁶ Let $V_{t+1}^{e,1}(\hat{\theta}; \psi, \varphi, m)$ denote firm value at this decision, when firm age and initial-product tenure equal one. Suppressing common states $(1, \psi, \varphi, m)$ on the right-hand side,

$$V_{t+1}^{e,1}(\hat{\theta}; \psi, \varphi, m) = \max \left\{ V_{t+1}^{ns}(\hat{\theta}, 1) + (n(\varphi) - 1)V_{t+1}^{ds}(\emptyset), \right. \\ \left. V_{t+1}^{cs}(\hat{\theta}, 1) + (n(\varphi) - 1)V_{t+1}^{ds}(\emptyset), V_{t+1}^{ds}(\emptyset) \right\}. \quad (2.20)$$

Thus, the initial retention decision accounts for the search opportunities preserved in the other lines.

Let $\tilde{V}_t^e(\varphi)$ denote operating value after φ is observed but before m is revealed. Then

$$\tilde{V}_t^e(\varphi) = \beta \bar{\pi}_{it+1}(\hat{\theta}_{ip_1t}^0, 0; 0, \varphi) + \beta(1 - \delta) \mathbb{E}_{m \mid \varphi} \left[\mathbb{E}_{\hat{\theta}' \mid (\hat{\theta}_{ip_1t}^0, 0), m} V_{t+1}^{e,1}(\hat{\theta}'; \psi_{it+1}, \varphi, m) \right]. \quad (2.21)$$

²⁴Pre-entry investments precede the realization of φ_i . Prototyping or production planning reveals productive potential, allowing startups to abandon before production after an unfavorable realization.

²⁵Uniform assignments and symmetric product markets allow us to suppress expectations over initial product assignments.

²⁶Upon endogenous exit at age one, the initial line receives $V_{t+1}^{ds}(\emptyset)$, while initially vacant lines receive zero. From age two onward, the payoff-equivalence assumption in Section 2.3.7 applies to all lines.

The startup operates if and only if its operating value is nonnegative. Assume single crossing in φ : once operation is optimal, it remains optimal at higher types. Then

$$V_t^e(\varphi) = \max\{\tilde{V}_t^e(\varphi), 0\}, \quad (2.22)$$

$$\underline{\varphi}_t^e = \inf\{\varphi : \tilde{V}_t^e(\varphi) \geq 0\}. \quad (2.23)$$

At a continuous interior crossing, $\tilde{V}_t^e(\underline{\varphi}_t^e) = 0$.²⁷

Before entry, potential startups choose (a_t^m, a_t^z) to maximize expected net entry value:

$$(a_t^m, a_t^z) \in \arg \max_{a^m, a^z > 0} \{\mathbb{E}_\varphi[V_t^e(\varphi)] - w_t f_t^e(a^m, a^z)\}. \quad (2.24)$$

We suppress the resulting dependence of $V_t^e(\varphi)$ on (m_{0t}, z_{0t}) . Free entry implies

$$\mathbb{E}_\varphi[V_t^e(\varphi)] = w_t f_t^e, \quad (2.25)$$

where $f^e \equiv f_t^e(a_t^m, a_t^z)$ is the startup resource cost at the optimum. In stationary equilibrium, $a^m = m_0$ and $a^z = z_0$. The mass of entrants is determined in equilibrium.

2.5 Equilibrium

2.5.1 Labor Market Clearing

Let $\xi_{ipt} \equiv (\hat{\theta}_{ipt}, \tau_{ipt}; a_{it}, \psi_{it}, \varphi_i, m_i) \in \mathcal{X}$ denote the state of an active product line after period- t search and entry, where $\mathcal{X}$ is the active product-line state space. From (2.7), its variable production labor demand in $t + 1$ is

$$\ell_{ipt+1}(\xi_{ipt}) = \left(\frac{\eta - 1}{\eta}\right)^\eta P_{pt+1}^\eta C_{pt+1} [\mathcal{B}(\xi_{ipt})]^\eta z(\varphi_i; z_0)^{\eta-1} w_{t+1}^{-\eta}. \quad (2.26)$$

²⁷The cutoff-based distribution formulas assume continuous H_φ , as in the quantitative specification.

Let μ_{pt} denote the measure of active product lines over $\mathcal{X}$ in market p at the end of period t . Aggregating across these lines gives total production labor demand:

$$L_{pt+1} = \left(\frac{\eta - 1}{\eta} \right)^\eta P_{pt+1}^\eta C_{pt+1} w_{t+1}^{-\eta} \int_{\mathcal{X}} \mathcal{B}(\xi)^\eta z(\varphi; z_0)^{\eta-1} \mu_{pt}(d\xi). \quad (2.27)$$

The labor market clearing condition is

$$\bar{L} = \sum_{p=1}^N L_{pt+1} + \mathcal{F}_{t+1}^o + \mathcal{F}_{t+1}^s + \mathcal{F}_{t+1}^e, \quad (2.28)$$

where $\mathcal{F}_{t+1}^o$, $\mathcal{F}_{t+1}^s$, and $\mathcal{F}_{t+1}^e$ denote labor used for operating fixed costs, search costs, and startup costs (including capability investments), respectively (Section 2.5.3).

2.5.2 Firm and Product-Line Distributions

Let $\xi_{it}^f \equiv (\Theta_{it}, a_{it}, \psi_{it}, \varphi_i, m_i) \in \mathcal{X}^f$ denote the firm-level state.²⁸ Let Γ_t denote the cross-sectional measure of firms over $\mathcal{X}^f$ after period- t search and entry.

Given equilibrium prices and policies, $\mathcal{Q}_{t+1}^f(A \mid \xi^f)$ is the probability that a firm in state ξ^f survives to the end of period $t + 1$ with its state in $A \subseteq \mathcal{X}^f$. This transition incorporates Bayesian updating, product-drop and search decisions, search outcomes and destination assignments, and changes in age and information-processing ability. The remaining probability corresponds to exogenous death or endogenous exit.

The distribution of startup states conditional on operating is

$$\Gamma_t^e(A) \equiv \frac{1}{N[1 - H_\varphi(\underline{\varphi}_t^e)]} \sum_{p=1}^N \int_{\varphi \geq \underline{\varphi}_t^e} \int \mathbf{1}_{\{\xi_{it}^{f,e}(p, \varphi, m) \in A\}} H_m(dm \mid \varphi; m_{0t}) H_\varphi(d\varphi), \quad (2.29)$$

where $\xi_{it}^{f,e}(p, \varphi, m)$ denotes an age-zero firm with one product assigned to market p at the common prior and $n(\varphi) - 1$ vacant lines. The firm measure evolves according to

$$\Gamma_{t+1}(A) = \int_{\mathcal{X}^f} \mathcal{Q}_{t+1}^f(A \mid \xi^f) \Gamma_t(d\xi^f) + \mathcal{M}_{t+1}^e \Gamma_{t+1}^e(A), \quad (2.30)$$

²⁸ Θ_{it} also records product assignments and, for the search exclusion set, the most recent destination associated with each vacant line. Initially vacant startup lines have no previous destination.

where $\mathcal{M}_{t+1}^e$ is the mass of operating entrants.

The active product-line measure μ_{pt} used in market clearing is obtained by projecting Γ_t onto active lines in market p . For $A \subseteq \mathcal{X}$,

$$\mu_{pt}(A) = \int_{\mathcal{X}^f} \sum_{k=1}^{n(\varphi)} \mathbf{1}_{\{k \text{ active}, p_k=p\}} \mathbf{1}_{\{\xi_k \in A\}} \Gamma_t(d\xi^f), \quad (2.31)$$

where p_k and ξ_k denote line k 's product assignment and state within ξ^f . Thus, μ_{pt} captures the mass and state distribution of active product lines in market p .

In stationary equilibrium, the firm measure satisfies

$$\Gamma(A) = \int_{\mathcal{X}^f} \mathcal{Q}^f(A \mid \xi^f) \Gamma(d\xi^f) + \mathcal{M}^e \Gamma^e(A), \quad (2.32)$$

with μ_p given by (2.31).²⁹

2.5.3 Aggregate Variables

Since each active product line uses f^o units of labor for its operating fixed cost, total labor used for operating fixed costs in period t is

$$\mathcal{F}_t^o = f^o \sum_{p=1}^N \mu_{p,t-1}(\mathcal{X}). \quad (2.33)$$

Let $\mathcal{S}_t(\xi^f)$ denote the expected number of product-line searches by a firm starting period t in state ξ^f , accounting for within-period uncertainty, exogenous death, and endogenous exit before search. Total labor used for search is

$$\mathcal{F}_t^s = f^s \int_{\mathcal{X}^f} \mathcal{S}_t(\xi^f) \Gamma_{t-1}(d\xi^f). \quad (2.34)$$

Online Appendix [OA.A.6](#) provides the expression using product-line measures.

Because only a fraction $1 - H_\varphi(\varphi_t^e)$ of startups operate, the gross mass paying pre-entry

²⁹Tracking portfolio states becomes computationally infeasible as scope increases. We compute equilibrium using product-line measures, including vacancies, and approximate endogenous exit probabilities from these marginals. Online Appendix [OA.A.5](#) provides the laws of motion and approximation.

costs is $\mathcal{M}_t^e/[1 - H_\varphi(\underline{\varphi}_t^e)]$. Total labor used for startup costs is therefore

$$\mathcal{F}_t^e = f_t^e(a_t^m, a_t^z) \frac{\mathcal{M}_t^e}{1 - H_\varphi(\underline{\varphi}_t^e)}. \quad (2.35)$$

Substituting the market-clearing price in (2.9) into the product-level price index gives

$$P_{pt} = \frac{\eta}{\eta - 1} w_t \left[\int_{\mathcal{X}} \mathcal{B}(\xi)^\eta z(\varphi; z_0)^{\eta-1} \mu_{p,t-1}(d\xi) \right]^{-\frac{1}{\eta-1}}. \quad (2.36)$$

Appendix A.4 provides the derivation.

In symmetric equilibrium, production labor demand is identical across product markets, so (2.28) implies $L_{pt} = (\bar{L} - \mathcal{F}_t^o - \mathcal{F}_t^s - \mathcal{F}_t^e)/N$. Combining this with (2.27) and (2.36) gives

$$C_{pt} = \frac{\bar{L} - \mathcal{F}_t^o - \mathcal{F}_t^s - \mathcal{F}_t^e}{N} \left[\int_{\mathcal{X}} \mathcal{B}(\xi)^\eta z(\varphi; z_0)^{\eta-1} \mu_{p,t-1}(d\xi) \right]^{\frac{1}{\eta-1}}. \quad (2.37)$$

Aggregate consumption is then $\bar{C}_t = N^{\nu/(\nu-1)} C_{pt}$.

2.5.4 Symmetric Stationary Recursive Competitive Equilibrium

Definition 1. A symmetric stationary recursive competitive equilibrium is a collection

$$\left\{ C, P, w, \mathcal{M}^e, a^m, a^z, m_0, z_0, y, p, \ell, L, \mathcal{F}^o, \mathcal{F}^s, \mathcal{F}^e, \right. \\ \left. V, V^{ns}, V^{cs}, V^{ds}, V^{e,1}, \tilde{V}^e, V^e, \mathcal{B}^{cs}, \mathcal{B}^{ds}, \underline{\varphi}^e, \Gamma, \Gamma^e, \mu \right\}$$

such that: (i) posterior beliefs follow (2.3) and (2.4), with $\mathcal{B}$ given by (2.8); (ii) given C, P, w , and beliefs, output y , prices p , and variable production labor demand ℓ satisfy (2.7), (2.9), and (2.26), respectively; (iii) product-line values satisfy (2.12), (2.13), (2.14), and (2.15), and search and drop thresholds satisfy (2.16) and (2.17); (iv) startup values, the operating cutoff, and capability investments satisfy (2.20), (2.21), (2.22), (2.23), and (2.24), with free entry given by (2.25); (v) capabilities satisfy (2.19), with $M = m_0$, $Z = z_0$, and hence $a^m = m_0$ and $a^z = z_0$; (vi) Γ^e is given by (2.29), Γ is invariant under (2.32), and μ is its product-line projection in (2.31); (vii) $L, \mathcal{F}^o, \mathcal{F}^s$, and $\mathcal{F}^e$ satisfy (2.27), (2.33), (2.34), and (2.35), and the labor market clears according to (2.28); and (viii) P and C satisfy (2.36) and (2.37),

with symmetry requiring $P_p = P$, $C_p = C$, $L_p = L$, and $\mu_p = \mu$ for all p .

Because only relative prices matter, we normalize the stationary wage to one, $w = 1$.

2.6 Model Implications

2.6.1 Output Wedge under Imperfect Learning

We compare production with the benchmark in which permanent expertise is known, holding aggregate prices and quantities fixed. Shrinkage toward the prior lowers expected output for sufficiently high-expertise products and raises it for sufficiently low-expertise products. Output is nonlinear in beliefs, so the expertise cutoff for expected underproduction differs from the prior mean.

Proposition 3 (Learning Frictions and Output Wedges). *Let y_{ipt}^{FI} denote optimal output when θ_{ip} is known but the contemporaneous demand shock remains unobserved. Define the output wedge by*

$$1 - \omega_{ipt} \equiv \frac{y_{ipt}^*}{y_{ipt}^{FI}} = \frac{\left[\mathcal{B}_{it-1}(\hat{\theta}_{ipt-1}, \tau_{ipt-1}) \right]^\eta}{\theta_{ip}^{\eta-1} \exp \left\{ \frac{(\eta-1)^2}{2\eta} \sigma_\varepsilon^2 \right\}}. \quad (2.38)$$

For $\hat{\sigma}_{ipt-1}^2 > 0$, the expected wedge conditional on true expertise is positive if and only if

$$\log \theta_{ip} > \theta_0 + \frac{\eta-1}{2} \left[\left(1 + \frac{1}{\eta} \right) \sigma_0^2 - \hat{\sigma}_{ipt-1}^2 \right], \quad (2.39)$$

and negative when the inequality is reversed.³⁰ Lower cumulative processing raises posterior variance and lowers this cutoff, expanding the range of expertise associated with expected underproduction and increasing its magnitude wherever it occurs.

Proof. See Appendix A.5. □

These comparisons isolate the intensive-margin effects of learning. Average wedges among continuing products additionally reflect endogenous selection.

³⁰Expectations average over signal histories at fixed tenure and processing path, without conditioning on product retention.

2.6.2 Product Switching under Imperfect Learning

Firms retain uncertain products partly for the value of further learning (Lemma A.1 in Appendix A.6). Information processing affects switching through two channels: the precision of accumulated information and the informativeness of future signals. We isolate the accumulated-information channel by raising current processing, ψ_{it} , while holding processing at all other dates fixed. This *fixed future processing benchmark* shows how unresolved uncertainty affects product selection while keeping replacement opportunities unchanged.³¹

The productivity and tenure comparisons use the baseline model and hold its age-dependent processing path fixed. For productivity, we impose two conditions on expected discounted fixed operating and search costs, including costs after future product replacements. At the drop threshold, retaining the existing product without search costs more than starting next period with a vacant line and following optimal decisions thereafter. At the search threshold, the existing product's extra cost relative to a new product is smaller than the search cost per successful replacement. Appendix A.7 states the conditions formally and explains their economic meaning. We maintain the threshold conditions of Proposition 1 and hold aggregate prices and quantities fixed. Comparisons concern feasible states; derivatives apply wherever they exist.

Proposition 4 (Information Processing and Product Switching).

- (i) *Information processing.* In the fixed future processing benchmark, higher ψ_{it} raises both thresholds for $\tau \geq 1$: $\partial \mathcal{B}_{it}^{ds}(\tau) / \partial \psi_{it} > 0$ and $\partial \mathcal{B}_{it}^{cs}(\tau) / \partial \psi_{it} > 0$.
- (ii) *Productivity.* Under conditions (A.8)–(A.9), higher z_i lowers the drop threshold, $\partial \mathcal{B}_{it}^{ds}(\tau) / \partial z_i < 0$, and raises the search threshold, $\partial \mathcal{B}_{it}^{cs}(\tau) / \partial z_i > 0$.
- (iii) *Product tenure.* At the same firm age and processing path, longer product tenure raises both thresholds: $\mathcal{B}_{it}^j(\tau + 1) > \mathcal{B}_{it}^j(\tau)$ for $j \in \{ds, cs\}$.

Proof. See Appendix A.7. □

At a given demand term, higher current processing and longer tenure leave less to learn about an existing product, which lowers the value of further experimentation and raises

³¹The fixed future processing path may rise with firm age.

both thresholds. Holding the current posterior fixed, higher constant processing from tomorrow onward has two opposing effects: better learning about the existing product lowers both thresholds, while improved replacement opportunities raise both, each holding the other fixed. The net effect depends on which dominates. Online Appendix [OA.A.10](#) establishes this decomposition and characterizes combined changes in both channels.

Under the conditions in part (ii), higher productivity makes weaker products worth keeping and search worthwhile at higher demand, widening the continue-and-search region. Processing and tenure raise both thresholds, with ambiguous effects on its width. Observed switching also depends on the demand distribution, which we examine quantitatively.

2.6.3 Returns to Pre-Entry Investment

Holding aggregate prices, quantities, and the normalized-type distribution fixed, higher managerial capability weakly raises expected entry value, while higher productivity capability raises it strictly if production occurs with positive probability.³² Firms can disregard additional information, so improved processing cannot reduce value. Higher productivity raises gross operating payoffs without directly changing operating or search fixed costs.

The capabilities are locally complementary if and only if better processing weakly raises expected discounted gross operating payoffs. A positive return to managerial investment alone does not suffice because its gains may come from fixed-cost savings.

3 Empirical Analysis

The model highlights how imperfect learning shapes product selection and firm growth. We turn to firm-product-level data from U.S. manufacturing to examine evidence for this mechanism. We estimate firms' product-specific expertise and study how product-switching decisions vary with expertise, firm age, and information processing. We examine the relationship between product switching and firm growth and performance.

³²Appendix [A.8](#) states the normalization and regularity conditions.

3.1 Data and Measurement

Our main data source is the U.S. Census Bureau’s quinquennial Census of Manufactures (CMF), linked to the Longitudinal Business Database (LBD). Our sample covers eight census years from 1977 to 2012. The CMF covers U.S. manufacturing establishments with at least one paid employee and reports product codes—five-digit SIC or seven-digit NAICS codes—and the value of shipments by product. We link CMF establishments to the LBD and use firm identifiers to aggregate establishment-level observations to the firm level. Online Appendix [OA.D.1](#) describes the data sources and construction.

We also use retrospective responses for 2005 from the 2010 Management and Organizational Practices Survey (MOPS), an establishment-level survey of U.S. manufacturing. We use the structured management score of [Bloom et al. \(2019\)](#), data availability (Q27), and data use intensity (Q28). We aggregate these measures to firms using value-of-shipments weights and link them to product-dropping outcomes over 2007–2012.

3.1.1 Product Definition and Characteristics

We define products at the five-digit SIC level, using 1987 SIC codes throughout the sample. For 1977–1997, we use product codes that the Census Bureau concords to 1987 SIC codes. For subsequent census years, we construct concordances using the 1997 CMF and successive Census bridge files, as detailed in Online Appendix [OA.D.1](#).

The CMF reports the value of shipments by product and, for some, physical quantities. We require positive shipment values and exclude administrative-record establishments, which need not report product-level shipments. We keep observations in the Census Bureau’s “tabbed sample,” used for official tabulations. We deflate shipment values to 1997 dollars using NBER-CES industry-level deflators ([Becker et al., 2013](#)). We measure firm-product tenure as the number of years the firm is observed producing it, supplementing the CMF with firm-product records from the Annual Survey of Manufactures (ASM).

We identify product additions and drops between consecutive census years. A firm adds a product between t and $t + 5$ if it does not produce the product in t but produces it in $t + 5$. A firm drops a product if it produces the product in t but not in $t + 5$. We construct these

indicators for census years t from 1982 to 2007.³³ We distinguish initial additions from re-entries and temporary drops from permanent drops. An initial addition is a product the firm has not previously been observed producing; a permanent drop is a product that does not reappear in the firm’s subsequent observed product portfolio. Our baseline analysis focuses on initial additions and permanent drops.

3.1.2 Firm Characteristics and Identifiers

We use firm age constructed by the Census Bureau following [Haltiwanger et al. \(2013\)](#). When a firm first appears, its age equals that of its oldest establishment; firm age increases annually thereafter. We define young firms as those aged five or less. We use the 1972 CMF to adjust left-censored ages and reduce misclassification of young firms. We measure firm size by total employment or value of shipments. Our baseline analysis uses the Census Bureau’s firm identifier, `lbd fid`; Online Appendix [OA.D.1](#) describes the age adjustment and alternative longitudinal firm identifiers constructed following [Dent et al. \(2018\)](#). All nominal values are deflated to 1997 dollars using industry-level deflators for each variable.

Summary Statistics Table [OA.9](#) in the Online Appendix reports means, standard deviations, and observation counts. The firm-level sample contains approximately 741,000 firm-year observations for firms with estimable product-specific expertise, requiring observations in at least two census years. Firms aged five or less account for 13% of these observations. Young firms are substantially smaller, employing 32.5 workers on average compared with 131.3 among mature firms. The table summarizes the firm-product characteristics and management measures used in the empirical analysis.

3.2 Empirical Results

Using these data, we examine how firms’ product-switching decisions vary with product-specific expertise, firm age, and information processing, and how these decisions relate to firm growth and performance.

³³We use the 1977 CMF to help identify prior production histories and measure firm-product tenure. The 2012 CMF provides the endpoint for the final five-year interval.

Table 1: Prevalence of Product Switching and Multi-Product Firms

	All firms		Young firms		Mature firms	
	Firms	Output	Firms	Output	Firms	Output
<i>Panel A: Product switching among continuing firms</i>						
Any switching	0.646	0.910	0.667	0.694	0.642	0.913
Add only	0.128	0.067	0.130	0.139	0.128	0.066
Drop only	0.182	0.169	0.143	0.146	0.187	0.169
Add and drop	0.336	0.674	0.394	0.409	0.327	0.678
<i>Panel B: Multi-product firms in the full sample</i>						
Multi-product firms	0.450	0.891	0.432	0.615	0.452	0.896

Notes: Entries are firm and output shares within each age group. Young firms are aged five or less; mature firms are aged six or more. Products are defined by five-digit SIC codes. Panel A covers continuing firms and measures additions and drops over five-year intervals; any switching is the sum of the three mutually exclusive categories. Panel B includes entrants and exiters. Online Appendix Tables [OA.10](#) and [OA.11](#) report the detailed results, including alternative samples and product definitions.

3.2.1 Prevalence of Product Switching

Panel A of Table 1 reports shares of firms engaging in product switching and their output shares by age group. To distinguish switching within continuing firms from changes associated with entry and exit, we focus on continuing firms. In the pooled sample, 65% of continuing firms switch products over a five-year interval, accounting for 91% of output among continuing firms. These patterns underscore the prevalence and economic importance of product switching, consistent with [Bernard et al. \(2010\)](#).

Product switching often involves both adding and dropping products, which we refer to as product churning. Approximately 52% of product-switching firms (34% out of 65%) both add and drop products; the remainder either add or drop products. Product churning is more prevalent among young firms, whereas dropping products without adding others is less common than among mature firms.

Panel B reports the prevalence of multi-product firms under our baseline five-digit definition. Multi-product firms constitute 45% of firms but account for 89% of output. They are prevalent even among young firms, accounting for 43% of firms and 62% of output in this age group. Online Appendix [OA.D.3](#) reports detailed statistics, including diversification across industries and sectors. Young continuing firms produce fewer products than mature

firms (1.6 versus 2.3), but add more (1.6 versus 1.2) and drop fewer (0.8 versus 1.0), as shown in Online Appendix Table [OA.12](#).

3.2.2 Product-Specific Expertise Estimation

We estimate firms' product-specific expertise to examine its relationship with product selection and resource reallocation. The CMF reports the value of shipments at the firm-product level. Shipment values, however, reflect both expertise and other determinants of product performance, including product life cycle dynamics ([Argente et al., 2024](#)), customer capital accumulation, and transitory demand shocks.

Our empirical specification is motivated by equation (2.11), which relates log real shipments to product-specific expertise, firms' beliefs and productivity, and product-market conditions. We capture persistent differences in expertise through firm-product fixed effects, in the spirit of [Bonhomme et al. \(2019\)](#), and estimate:

$$y_{ipjt} = \theta_{ip} + X_{ipjt}\gamma_1 + X_{ijt}\gamma_2 + \delta_{pt} + \varepsilon_{ipjt}, \quad (3.40)$$

where y_{ipjt} denotes log real shipments of product p by firm i in industry j and census year t . The firm-product fixed effect θ_{ip} captures persistent differences in shipment values conditional on controls and fixed effects. We interpret this component as an empirical measure of product-specific expertise, reflecting how well a firm is suited to a product. Appendix [B.2](#) discusses its mapping to expertise in the model.

X_{ipjt} includes firm-product tenure and its square to account for product life cycle dynamics within firms, including changes in shipments as firms gain production experience and build a customer base. Firm-level controls, X_{ijt} , account for differences in scale, resources, and production inputs that affect shipments across a firm's products.³⁴ Product scope and total assets also help account for alternative explanations for product-switching patterns over the firm life cycle, as discussed below. Product-year fixed effects, δ_{pt} , absorb demand

³⁴These controls include log employment, firm age, total assets, number of products, number of establishments, advertising and software expenditures, skill and capital intensity, and productivity. Skill and capital intensity are measured by the log ratios of nonproduction employment and capital to total employment, respectively. We also include the logs of production employment, equipment and structures capital stocks, material costs, and energy costs.

and technology conditions common to firms producing the same product in a year, including persistent differences across product markets.

We estimate θ_{ip} using the 1977–2012 CMF, supplemented with ASM observations from 1976–2016. For ease of interpretation, we rescale the estimated fixed effects to the unit interval as $q_{ip} \equiv (\hat{\theta}_{ip} - \underline{\theta})/(\bar{\theta} - \underline{\theta})$, where $\underline{\theta}$ and $\bar{\theta}$ denote the minimum and maximum estimated fixed effects in the sample. This transformation preserves their ordering. Section 3.3 examines sensitivity to estimation windows, specifications, normalization methods, and statistical inference.

3.2.3 Product Selection

The model links product selection to firms' information about their product-specific expertise. Using our expertise measure, we examine which products firms retain, drop, and add, and how these decisions vary with firm age, management practices, and data availability, which shape information-processing precision in the model.

Selection on Product Performance We first examine whether firms are more likely to drop products with lower contemporaneous shipments or estimated expertise. This establishes the overall pattern of selection before examining how it varies with firm age and information processing. Let $\mathcal{I}_{ipjt}^{drop}$ equal one if firm i in industry j permanently drops product p between census years t and $t + 5$. Our specifications are nested in:

$$\mathcal{I}_{ipjt}^{drop} = \alpha + \beta_1 Q_{ipt} + \beta_2 \mathcal{I}_{ijt}^{young} + \beta_3 Q_{ipt} \times \mathcal{I}_{ijt}^{young} + X_{ipjt} \gamma_1 + X_{ijt} \gamma_2 + \delta_i + \delta_{pt} + \varepsilon_{ipjt}. \quad (3.41)$$

Here, Q_{ipt} denotes either log real shipments or normalized product-specific expertise, q_{ip} , and $\mathcal{I}_{ijt}^{young}$ equals one if firm i is aged five or less in year t . The vectors X_{ipjt} and X_{ijt} contain the firm-product and firm-level controls used in expertise estimation, respectively. When $Q_{ipt} = q_{ip}$, we also control for log real shipments to distinguish selection on expertise from selection on contemporaneous performance. Firm fixed effects, δ_i , absorb persistent differences across firms, while product-year fixed effects, δ_{pt} , absorb conditions common to

Table 2: Probability of Product Drop

	$Drop_{t:t+5}$	$Drop_{t:t+5}$	$Drop_{t:t+5}$	$Drop_{t:t+5}$	$Drop_{t:t+5}$
$\log(PV_{pit})$	-0.052*** (0.002)			-0.052*** (0.002)	
q_{ip}		-1.074*** (0.021)			-1.081*** (0.022)
Young Firm $_{ijt}$			-0.075*** (0.003)	-0.141*** (0.007)	-0.368*** (0.022)
$\log(PV_{pit}) \times \text{Young}_{ijt}$				0.010*** (0.001)	
$q_{ip} \times \text{Young}_{ijt}$					0.497*** (0.043)
Observations	1,622,000	1,622,000	1,622,000	1,622,000	1,622,000
Fixed effects	pt, i	pt, i	pt, i	pt, i	pt, i
Controls	Full	Full	Full	Full	Full

Notes: Estimates for product-year, firm fixed effects, and the constant are suppressed. Robust standard errors clustered at the five-digit SIC product level are reported in parentheses. Observations are unweighted, and observation counts are rounded due to the Census Bureau disclosure avoidance procedures. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

firms operating in the same product market.

The first two columns of Table 2 report estimates with all controls, omitting the young-firm indicator and its interaction with product performance.³⁵ Column 1 shows that firms are more likely to drop products with lower contemporaneous shipments. Column 2 shows that firms are also more likely to drop products with lower expertise, even conditional on contemporaneous shipments. A one-standard-deviation increase in expertise is associated with a 26.2-percentage-point decline in the probability of dropping a product over the next five years. These findings indicate that product dropping is systematically related to both current performance and persistent firm-product characteristics, consistent with selection toward high-expertise products.

Product Dropping over the Firm Life cycle Lemma 1 shows that lower information-processing ability strengthens the shrinkage of beliefs toward the prior mean. Proposition 4(i) isolates a related selection channel: with future processing fixed, less precise

³⁵Full regression results appear in Online Appendix OA.D.7.

accumulated information lowers both thresholds. Firm age also changes future learning opportunities, so we examine empirically how product dropping and its relationship with expertise differ between young and mature firms. We first estimate equation (3.41) with the young-firm indicator, omitting product performance and the interaction term. Column 3 of Table 2 shows that young firms are 7.5 percentage points less likely to drop a product over the next five years.³⁶

Two controls are particularly relevant to interpreting age differences in product dropping. Product scope accounts for differences in firms' portfolios: dropping a firm's sole product entails exit unless the firm adds another product. Total assets help account for financial constraints through the role of physical assets as collateral, since financing difficulties may limit firms' ability to replace existing products. Together, these controls help address alternative explanations for age differences in product dropping.

To examine whether selection on product performance also varies with firm age, we estimate the unrestricted specification in equation (3.41), interacting each performance measure with the young-firm indicator. Here, β_3 captures the difference between young and mature firms in the relationship between product performance and dropping. Columns 4 and 5 report positive interaction coefficients for both contemporaneous shipments and product-specific expertise. Thus, product dropping is less sensitive to performance among young firms. The estimates in column 5 imply that a one-standard-deviation increase in expertise is associated with a 26.4-percentage-point decline in the dropping probability among mature firms, compared with 14.2 percentage points among young firms. Selection on expertise is therefore approximately 46% weaker among young firms.

These patterns are consistent with the model's learning mechanism, in which less precise information among young firms weakens selection toward high-expertise products. The results also hold when we replace the young-firm indicator with continuous firm age, as reported in Online Appendix Table OA.29.

The product addition results provide complementary evidence. Young firms are 14.8 percentage points more likely to add a product over the next five years. In the sample used to assess addition quality, however, they are 11.2 percentage points less likely to have additions

³⁶Full regression results with and without firm-level controls appear in Online Appendix OA.D.7.

Table 3: Probability of Product Drop

	$Drop_{t:t+5}$	$Drop_{t:t+5}$	$Drop_{t:t+5}$
q_{ip}	-1.055*** (0.041)	-1.127*** (0.033)	-1.119*** (0.035)
$q_{ip} \times m_{ijt}$	-1.036*** (0.082)	-0.832*** (0.062)	-0.912*** (0.068)
Observations	1,622,000	1,622,000	1,622,000
Fixed effects	pt, i	pt, i	pt, i
Controls	Full	Full	Full
Measure	Structured management	Data availability	Data use intensity

Notes: q_{ip} denotes normalized product-specific expertise, and m_{ijt} denotes the measure indicated in each column. Column (1) uses the structured management score of Bloom et al. (2019), column (2) uses data availability (Q27), and column (3) uses data use intensity (Q28). Measures refer to 2005, reported in the 2010 MOPS, and enter interacted with the 2007 indicator. Missing measures are set to zero; a missing-measure indicator and its interaction with q_{ip} are included. Management main effects, missingness controls, fixed effects, and the constant are not reported. Standard errors clustered at the five-digit SIC product level appear in parentheses. Observations are unweighted, and observation counts are rounded under Census Bureau disclosure avoidance procedures. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

whose average expertise exceeds that of their existing portfolio. Online Appendix OA.D.4 presents the specifications and results. Together, the dropping and addition results suggest that young firms experiment with products but exhibit weaker selection on expertise, consistent with learning frictions early in the firm life cycle.

Product Dropping and Information Processing In the model, management and data availability shape information-processing precision and thereby firms' beliefs about product-specific expertise. We examine whether product selection varies with management practices, data availability, and data use. In equation (3.41), we set $Q_{ipt} = q_{ip}$ and replace the young-firm indicator with m_{ijt} , the corresponding 2005 measure interacted with the 2007 indicator. Missing measures are set to zero, and each specification includes a missing-measure indicator and its interaction with q_{ip} . Online Appendix OA.D.7 presents the specification and full results.

Table 3 reports negative and statistically significant interaction coefficients across all three measures. Thus, the dropping probability declines more sharply with expertise among firms with higher management scores, greater data availability, or more intensive data use. These results complement the firm-age evidence and are consistent with more precise information

strengthening selection toward high-expertise products.

3.2.4 Product Switching and Firm Performance

Having documented how firms select products, we next examine how product additions and drops relate to firm growth and performance.

Online Appendix [OA.D.5](#) shows that product additions are positively associated with growth in real value added, employment, total factor productivity, and labor productivity, while product drops are negatively associated with growth. Under the additive specification, firms that add and drop products have higher predicted growth than firms that do neither. The implied differences are 0.076 log points for real value added, 0.071 for employment, and 0.049 for total factor productivity and labor productivity.

Decomposition of Firm-Level Growth in Output and Product Expertise To examine the role of within-firm reallocation through product switching, we conduct an accounting decomposition of firm-level growth. We measure growth as the change in the shipment-weighted average of log real product shipments and similarly examine changes in shipment-weighted product-specific expertise. The decomposition distinguishes within-product changes, shifts in output shares among continuing products, their interaction, and contributions from product additions and drops. Online Appendix [OA.D.6](#) provides the formula and detailed results.

Table 4 presents the normalized contributions, which we refer to as shares.³⁷ Panel A shows that growth in our shipment-based measure is primarily driven by within-product growth and new product additions. For young firms, however, new product additions play a more prominent role than within-product growth. The combined contribution of additions and drops is approximately 56% for young firms and 44% for mature firms, or 28% higher for young firms. Panel B shows that product additions also account for a larger share of the increase in shipment-weighted expertise among young firms.

³⁷Average five-year growth in the shipment-based measure is 0.112, 0.662, and 0.104 log points for all, young, and mature firms, respectively. The corresponding changes in the shipment-weighted average of θ_{ip} are 0.017, 0.130, and 0.016.

Table 4: Product Switching and Firm Growth

	All firms	Young firms	Mature firms
<i>Panel A: Growth in shipment-weighted log product shipments</i>			
Within-product growth	0.543	0.444	0.545
Between-product reallocation	0.008	0.002	0.009
Cross term	0.009	-0.005	0.010
Product additions	0.424	0.554	0.422
Product drops	0.015	0.005	0.015
<i>Panel B: Growth in shipment-weighted product-specific expertise</i>			
Between-product reallocation	0.318	0.126	0.320
Product additions	0.574	0.853	0.572
Product drops	0.107	0.022	0.108

Notes: Entries are normalized contributions to five-year changes in the shipment-weighted average of log real product shipments (Panel A) and estimated product-specific expertise, θ_{ip} (Panel B). Young firms are aged five or less; mature firms are aged six or more. The drop row includes the minus sign in Online Appendix equation (OA.D.3). Within and cross terms are zero for time-invariant expertise. Shares may not sum to one after rounding. Online Appendix OA.D.6 provides the decomposition and results.

Taken together, these findings connect product switching to firm performance. The greater role of product additions among young firms, alongside their weaker selection on expertise, is consistent with the model’s learning mechanism: firms experiment with products, while limited information-processing ability constrains selection toward high-expertise products early in the firm life cycle.

3.3 Robustness Tests

We assess whether age differences in product selection are sensitive to expertise estimation, normalization, statistical inference, and firm exit. Online Appendix OA.D.9 reports results. Observation histories may affect the precision of estimated expertise, particularly for young firms. We therefore re-estimate expertise using only the first two observations per firm-product pair and use 5-, 10-, and 15-year rolling estimation windows. The interaction between expertise and the young-firm indicator remains positive and statistically significant across these specifications, although its magnitude varies.

Expertise estimates may also capture firm-level determinants of shipments not fully accounted for by baseline controls. We vary these controls, including a specification without

them, for which the positive age interaction persists. The age difference in selection also persists under standardization and percentile rankings within industry-year or product-year cells. With bootstrapped standard errors, the expertise coefficient remains negative and the young-firm interaction positive, with both statistically significant. Finally, the main age-related selection pattern persists after excluding firms that subsequently exit. Online Appendix [OA.D.11](#) discusses alternative specifications with firm-year fixed effects and an AKM-style residual-based measure of expertise.

4 Quantitative Analysis

We quantify how learning about product-specific expertise shapes young firms' growth and exit. We extend the model to persistent demand shocks and calibrate it to U.S. manufacturing. After assessing model fit, we examine how information processing connects product selection to firm continuation and growth. We then evaluate a subsidy to startups' managerial investment costing 0.1% of baseline consumption expenditure. A paired entrant experiment identifies production, portfolio, and exit responses, while the general-equilibrium comparison captures adjustments in prices, entry, and resource allocation. Finally, a restricted-planner benchmark clarifies the role of capability investment and spillovers across cohorts.

4.1 Persistent Product-Market Shocks

For the quantitative analysis, we allow the transitory component of product-market demand to be persistent. We retain $\log \theta_{ipt} = \log \theta_{ip} + \varepsilon_{ipt}$ and assume

$$\varepsilon_{ip,t+1} \mid \varepsilon_{ipt} \sim N(\rho \varepsilon_{ipt}, (1 - \rho^2) \sigma_\varepsilon^2), \quad 0 \leq \rho < 1. \quad (4.42)$$

Stationary variance remains σ_ε^2 , and the additional signal (2.2) is unchanged. Persistence reduces the independent information in successive market outcomes: each observation after the first contributes precision $(1 - \rho)/[(1 + \rho)\sigma_\varepsilon^2]$ about permanent expertise.

Current realized demand also predicts future demand, so the product-line state addition-

ally retains $\log \theta_{ipt}$. For positive tenure, the demand term becomes

$$\begin{aligned} \log \mathcal{B}_{it}(\hat{\theta}_{ipt}, \log \theta_{ipt}, \tau_{ipt}) &= \frac{\eta - 1}{\eta} \left[\rho \log \theta_{ipt} + (1 - \rho) \hat{\theta}_{ipt} \right] \\ &\quad + \frac{1}{2} \left(\frac{\eta - 1}{\eta} \right)^2 \left[(1 - \rho)^2 \hat{\sigma}_{ipt}^2 + (1 - \rho^2) \sigma_\varepsilon^2 \right]. \end{aligned} \quad (4.43)$$

New products retain the common prior and initial demand term $\mathcal{B}^0$. Production and product-switching decisions use the revised demand term and Bayesian transition, with the same timing, entry, and exit conventions. Online Appendix [OA.B.1](#) provides the full equations. Setting $\rho = 0$ recovers the baseline model.

4.2 Parameterization

Among potential entrants, fundamental type φ_i and managerial ability m_i are jointly log-normal:

$$\begin{pmatrix} \log \varphi_i \\ \log m_i \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ \log m_0 \end{pmatrix}, \begin{pmatrix} \sigma_\varphi^2 & \rho_{m\varphi} \sigma_\varphi \sigma_m \\ \rho_{m\varphi} \sigma_\varphi \sigma_m & \sigma_m^2 \end{pmatrix} \right). \quad (4.44)$$

The parameters σ_φ and σ_m govern dispersion in log productivity and log managerial ability, while $\rho_{m\varphi}$ governs their correlation before endogenous selection.

Productivity and potential scope are

$$z_i = z_0 \varphi_i, \quad n_i = \max \{1, \lfloor n_0 \varphi_i^{\alpha_n} \rfloor\}. \quad (4.45)$$

Here $n_0 > 0$ scales potential scope and $\alpha_n \geq 0$ governs its responsiveness to firm type.³⁸ Realized scope depends on product-switching decisions.

We proxy accumulated data by

$$D_{it} = D_0 + n_i(a_{it} + 1), \quad a_{it} > 1. \quad (4.46)$$

At age one, $D_{it} = D_0 + a_{it} + 1$ reflects data from the initial product. Subsequently, potential scope proxies for the scale of data generation, while $D_0 \geq 0$ captures other data sources.

³⁸We choose α_n so that the maximum n_i on the computational grid is well below $N/2$.

The data path is determined by firm age and permanent type.

Information-processing ability combines managerial ability with a CES aggregator of experience and data:

$$\psi_{it} = m_i^{\gamma_m} \left\{ 1 + (\gamma_a + \gamma_D) \left[\frac{\gamma_a}{\gamma_a + \gamma_D} [\log(1 + a_{it})]^{\frac{\kappa-1}{\kappa}} + \frac{\gamma_D}{\gamma_a + \gamma_D} [\log D_{it}]^{\frac{\kappa-1}{\kappa}} \right]^{\frac{\kappa}{\kappa-1}} \right\}. \quad (4.47)$$

The parameters $\gamma_a, \gamma_D > 0$ govern the contributions of experience and data, and $\kappa > 0$ governs their substitutability. Managerial ability scales information processing with elasticity $\gamma_m > 0$. At age zero, $\psi_{it} = m_i^{\gamma_m}$. The logarithmic inputs allow the contributions of experience and data to increase at diminishing rates.

4.3 Solution Algorithm

The product-line state is $(\hat{\theta}, \log \theta, \tau, a, \varphi, m)$, with two time-varying continuous states, $(\hat{\theta}, \log \theta)$. We normalize the wage to one and initialize $m_0, z_0, P^\eta C$, and the value functions V^{ns}, V^{cs} , and V^{ds} . We then: (i) construct $z_i, m_i, n_i, D_{it}, \psi_{it}$, the posterior-variance schedule, the joint belief and demand transition, and $\mathcal{B}(\xi)$; (ii) solve the product-line Bellman equations and associated policies to convergence; (iii) compute the expected value of entry and the entry cutoff $\underline{\varphi}^e$; (iv) compute the stationary distribution of active and vacant product lines per unit mass of operating entrants, accounting for product switching, endogenous and exogenous exit, search outcomes, and entry; and (v) update $P^\eta C$ to satisfy free entry, repeating Steps (ii)–(iv) after each update.

Once free entry holds, we recover the entrant mass $\mathcal{M}^e$ from labor-market clearing and scale the stationary distribution. We then recover P and C consistent with the equilibrium $P^\eta C$. Finally, we update (m_0, z_0) using the pre-entry capability choice in (2.24), taking aggregate equilibrium objects and inherited capabilities as given. We repeat the equilibrium and capability updates until (m_0, z_0) and all equilibrium objects converge.

Table 5: Internally Calibrated Parameters and Target Moments

Par.	Description	Value	Target moment	Data	Model
σ_ε	Demand shock SD	0.890	SD log TVS	1.54	2.73
α_z^e	Productivity investment cost	3.47e−6	Mean log LP	3.80	3.69
σ_φ	SD log productivity	0.760	SD log LP	0.79	0.38
α_m^e	Management investment cost	3.17e−2	Mean $Q28$	0.43	0.62
σ_m	SD log management	0.250	SD $Q28$	0.25	0.30
$\rho_{m\varphi}$	Management–productivity corr.	0.300	Corr($Q28, \log z$)	0.17	0.16
n_0	Product scope scale	0.053	Mean product scope	2.19	2.20
α_n	Product scope curvature	2.371	SD product scope	2.89	1.90
γ_m	Management elasticity	8.000	Drop reg.: $q \times Q28$	−0.91	−0.62
γ_a	Experience weight	0.300	Drop reg.: $q \times Young$	0.50	0.17
γ_D	Data weight	0.200	Drop reg.: $q \times Q27$	−0.83	−0.48
f^o	Operating fixed cost	0.002	Share dropping	0.64	0.53
f^s	Search fixed cost	0.009	Share adding	0.34	0.53
x	Search success probability	0.335	No. products added	1.27	1.34
$\bar{f}^e$	Fixed entry cost	0.033	Startup rate (%)	7.48	7.60

Notes: SD denotes standard deviation; TVS, total value of shipments; and LP, labor productivity. Product scope is realized product scope. q is the normalized expertise measure and *Young* is the young-firm indicator used in the empirical product-drop regression. Dropping and adding shares refer to firms and are fractions; the annual startup rate is in percent. The implied capability levels and total startup resource cost are $m_0 = 0.940$, $z_0 = 7.300$. Values are rounded.

4.4 Calibration

We set 12 parameters using external evidence or benchmark values and jointly calibrate 15 parameters to match empirical moments in Table 5. One period represents one year. We normalize $\bar{L} = 1$, set $D_0 = 0$, and use $N = 1,400$ product markets, corresponding to five-digit SIC products. We set $\beta = 0.966$, based on an average ten-year real interest rate of 3.54% over 1982–2007.³⁹ The substitution elasticities are $\eta = 6.6$ within product markets and $\nu = 4.0$ across markets, following Broda and Weinstein (2006). We set $\delta = 0.039$ to the average annual exit rate of manufacturing firms aged 26 or older in the BDS over 2004–2007.⁴⁰ We set $\theta_0 = -0.010$ and $\sigma_0 = 0.2125$, using the mean and standard deviation of the firm-product fixed effects estimated from the CMF, each divided by $\eta - 1 =$

³⁹Federal Reserve Bank of Cleveland, [10-Year Real Interest Rate](#), retrieved from FRED.

⁴⁰We use this shorter period because firms with left-censored birth years can by then be unambiguously classified as age 26 or older.

5.6.⁴¹ The remaining predetermined parameters are demand-shock persistence $\rho = 0.8$, the experience–data substitution elasticity $\kappa = 1.5$, entry-investment curvature $\chi = 6$, and inter-cohort capability transferability $(\zeta_m, \zeta_z) = (0.15, 0.05)$.⁴²

Table 5 reports fifteen internally calibrated parameters and target moments computed for the CMF sample used in the empirical analysis. The parameters are calibrated jointly; the row associations indicate the principal moments informing each parameter. Following the MOPS scoring procedure, we discretize the model counterparts of data availability ($Q27$) and data-driven decision making ($Q28$) into five categories from zero to one in increments of 0.25. Product-switching moments use five-year intervals.

The value of shipments, productivity, and $Q28$ moments discipline demand and firm heterogeneity and capability-investment costs, while product-scope moments inform (n_0, α_n) . The interactions of normalized expertise q with $Q28$, young-firm status, and $Q27$ in the product-drop regression discipline $(\gamma_m, \gamma_a, \gamma_D)$. Product dropping, additions, and the startup rate inform $(f^o, f^s, x, \bar{f}^e)$.

4.5 Model Fit and External Validation

The model closely matches the startup rate and mean product scope and reproduces the signs of the product-dropping interactions (Table 5). It overstates sales dispersion and understates productivity dispersion and interaction magnitudes. The model captures the concentration of firms in small product portfolios and declining exit with age, while understating young-firm exit and the firm and employment shares at ages 26 and above (Figure 3).

Untargeted comparisons also support the importance of demand heterogeneity. Demand-related components account for 64.4% of the variance of log deflated sales, compared with the 75.8% firm-appeal contribution in Hottman et al. (2016). Entrants’ demand and productivity gaps relative to old continuing firms are -1.39 and -0.05 log points, respectively; the corresponding plant-level estimates in Foster et al. (2016) are -0.89 and 0.02. Thus, both comparisons indicate a substantially larger demand disadvantage.

⁴¹See Appendix B.2 for this normalization.

⁴²We set $\chi = 6 > \eta - 1 = 5.6$ so investment costs eventually dominate productivity returns. A common exponent for both capabilities saves parameters, with α_m^e and α_z^e capturing differences in investment costs.




Figure 3: Firm distributions: model and data

Notes: Product scope is the number of currently produced products. The age-26+ group includes left-censored firms in the data.

Applying the employment variance decomposition of [Sterk et al. \(2021\)](#) to simulated firms, we find that ex ante heterogeneity accounts for 67.8% of log-employment variance at age five, compared with 62.1% for manufacturing firms in their data.⁴³ Online Appendix [OA.C.1](#) reports comparisons, including within-firm sales concentration.

4.6 Selection Within and Across Firms

We now use the calibrated model to examine how information processing connects product selection within firms to firm continuation and growth. Proposition 4 isolates the accumulated-information channel under fixed future processing and characterizes productivity effects under the stated cost conditions. Here we evaluate selection numerically with the full age-dependent processing path and persistent demand shocks ($\rho = 0.8$).

Figure 4 shows that firms' highest product-specific expertise generally rises with age and is higher among firms with greater managerial ability or productivity. These profiles reflect portfolio adjustment within firms and selection among continuing firms. Figure 5 connects

⁴³The ex ante share declines with age in the data but remains high in the model.




Figure 4: Highest product-specific expertise over the firm life cycle

Notes: Unweighted average across continuing firms of each firm's maximum expertise among currently produced products. Groups are fixed at entry. Faint lines show unsmoothed means; solid lines apply centered five-age smoothing with shorter windows at the endpoints. Profiles reflect both portfolio adjustment and changes in the composition of continuing firms.




Figure 5: Continuation and large-firm status at age 20

Notes: Probabilities pool firm-age masses within fixed productivity–managerial-ability groups. Continuation is measured relative to original entrants; large-firm status is conditional on continuation. Large firms have employment at or above the age-specific 90th percentile among active firms. Age-20 statistics pool firm-age masses over ages 18–22.

these patterns to firm continuation and size. Among high-productivity firms at age 20, those with high managerial ability have a 4.61 percentage-point lower continuation probability but a 5.94 percentage-point higher probability of being large conditional on continuation than those with low managerial ability.^{44,45}

⁴⁴The joint probability of being active and large is also higher: 5.59% versus 4.39%.

⁴⁵Potential scope also links product selection to firm exit. Low- and middle-productivity firms are restricted to one product, whereas high-productivity firms can drop weak products while retaining others. Middle-productivity firms replace products more intensively than low-productivity firms, increasing their exposure

These selection patterns highlight the distinction between current output and productive potential. Comparing high- and low- φ_i firms with comparable output, the probability of being active and in the top output decile ten years later is higher by 10.14 percentage points at initial age five and 12.37 percentage points at initial age ten.⁴⁶ Current output therefore provides an incomplete measure of firms' growth prospects. Online Appendix [OA.C.2](#) describes the comparison and reports additional outcomes.

4.7 Policy Experiment: Improving Information Processing

We next quantify how improved information processing affects firm growth, selection within and across firms, and aggregate outcomes. We consider a subsidy to managerial investment for entrants above the entry cutoff, provided before initial production. The program uses resources equal to 0.1% of baseline consumption expenditure and raises m_0 from 0.94 to 1.23, holding z_0 and the entry cost fixed. The experiment improves firms' ability to learn without directly changing productivity or product-specific expertise. We first examine how baseline entrants respond at baseline aggregate prices, then account for equilibrium adjustments in prices, entry, and resource allocation.

Young Firm Selection and Growth. We first conduct a paired counterfactual experiment, simulating each baseline entrant with and without the subsidy while holding productive potential, initial product-specific expertise, and random draws fixed.⁴⁷ We decompose gains in the cohort's cumulative real sales into production adjustments, product portfolio adjustments, and firm continuation. We aggregate sales across the cohort, deflate them by the baseline product-market price index, and sum from age zero through the displayed age without discounting, reporting totals per original entrant. Sales are zero after exit, capturing changes in continuation without conditioning on survival.

Production adjustments measure the effect of improved information when quantities are to uncertain products. With no other products to sustain production, dropping a weak product triggers exit, contributing to their lower continuation probability.

⁴⁶In our model, φ_i determines both productivity and potential scope.

⁴⁷We exclude changes in entry and equilibrium resource reallocation, which we examine next.




Figure 6: Managerial investment and entrant outcomes

Notes: Both panels follow the same baseline entrants at baseline aggregate prices. The left reports continuation among original entrants; the right decomposes cumulative real-sales changes per original entrant. Contributions are percentage points of cumulative baseline real sales; diamonds show their sum. Both panels use centered five-year moving averages. Online Appendix [OA.C.3](#) describes cohort construction.

reoptimized but baseline product and exit decisions are held fixed.⁴⁸ We then allow product switching and exit to adjust. At each age, portfolio adjustments capture the sales effects of product switching among firms producing in both paths. Firm continuation measures sales gained when the subsidy delays exit, net of sales forgone when it accelerates exit. We cumulate these contributions over age and normalize them by cumulative baseline real sales, so they sum to the total percentage change.⁴⁹

Figure 6 shows that improved information processing raises cumulative sales despite lower firm continuation. Continuation at age 20 falls from 20.14% to 17.99%, while cumulative real sales rise by 2.83%. Production adjustments contribute 1.88 percentage points, portfolio adjustments contribute 1.09 percentage points, and firm continuation subtracts 0.14 percentage points. These components account for 66.4%, 38.6%, and -5.0% of the cumulative sales gain, respectively. The portfolio-adjustment contribution rises from 0.11 percentage points at age five to 1.09 at age 20. Improved information processing therefore raises cumulative sales through both production adjustments within existing portfolios and changes in the products firms operate.

To examine selection early in the life cycle, we group the same entrants by initial expertise

⁴⁸This captures the intensive-margin response associated with the output wedge in Proposition 3.

⁴⁹The decomposition is sequential: interactions between improved information and changes in product and exit decisions are assigned to the portfolio-adjustment and firm-continuation components.

Table 6: Initial product expertise and young firm outcomes (%)

Outcome around age five	Lower initial expertise		Higher initial expertise	
	Baseline	Higher m_0	Baseline	Higher m_0
Operating with initial product	19.49	14.11	38.36	37.85
Operating without initial product	20.09	17.09	20.26	17.82
Exited	60.42	68.81	41.38	44.32

Notes: Lower and higher initial expertise denote log expertise below and above the prior mean θ_0 . Groups are fixed across paths. Entries are percentages of original entrants within each group with potential scope $n_i \leq 4$, for whom the initial product can be tracked; these firms account for 94.33% of entrants. The three outcomes are mutually exclusive and exhaustive. Values average ages 3–7. Rounding may prevent columns from summing to 100.

and compare their outcomes around age five (Table 6). The probability of still operating the initial product falls by 5.39 percentage points among entrants with lower initial expertise, compared with 0.51 points among those with higher expertise. Exit rises by 8.39 and 2.95 points, respectively. Better information therefore changes young firms' retention and exit more strongly when they begin with weak products. These outcomes are measured relative to original entrants, so retention reflects both product choices and exit.

To assess whether portfolio adjustments and changes in continuation reflect better-informed selection, we compare firms' product-switching and exit decisions at each realized state with a full-information benchmark. The benchmark reveals product-specific expertise and reoptimizes current and future decisions.⁵⁰ The share of product-switching decisions that differ from this benchmark falls from 21.70% to 13.76% over ages 0–5 and from 18.22% to 10.72% over ages 0–20 (Table 7). The corresponding firm-exit shares fall from 12.32% to 7.78% and from 8.60% to 5.01%.

These results quantify how information processing affects firm growth through production adjustments and product selection. Improved information aligns switching and exit decisions more closely with the full-information benchmark. Gains from production and portfolio adjustments more than offset the net sales loss from changes in firm continuation, making continuation alone an incomplete measure of program benefits.

⁵⁰Future demand shocks and the expertise of undiscovered products remain unknown.

Table 7: Decisions differing from the full-information benchmark (%)

Decision	Ages 0–5		Ages 0–20	
	Baseline	Higher m_0	Baseline	Higher m_0
Product switching	21.70	13.76	18.22	10.72
Firm exit	12.32	7.78	8.60	5.01

Notes: Product-switching decisions are continue without search, continue and search, and drop and search. Decisions at each realized state are compared with the full-information benchmark at baseline aggregate prices. The benchmark reveals product-specific expertise and reoptimizes current and future decisions; future shocks and undiscovered expertise remain unknown.

Table 8: Aggregate effects of the managerial investment subsidy

<i>A. Equilibrium effects</i>	
Outcome	Change (%)
Consumption and aggregate productivity	0.54
Output	0.10
Firm mass	-7.67
<i>B. Decomposition of the consumption change</i>	
Component	Contribution (pp)
Firm mass	-8.01
Firm composition	6.26
Within-group performance	2.29
Total	0.54

Notes: Panel B reports an accounting decomposition in percentage points of baseline consumption. Groups are defined by firm age, φ , and baseline managerial ability. Within-group performance measures mean deflated sales per firm and includes changes in firm composition. Interaction terms are allocated symmetrically.

Selection and Aggregate Outcomes. The paired experiment isolates production, portfolio, and exit responses for a fixed entrant cohort at baseline aggregate prices. We now examine how these responses affect aggregate outcomes when prices, entry, and resource allocation adjust, comparing stationary equilibria under the same subsidy.

After accounting for the program’s resource cost, consumption rises by 0.54% and output by 0.10%, while firm mass falls by 7.67% (Table 8).⁵¹

To connect these aggregate outcomes to changes in selection and performance, we decompose the consumption gain into firm mass, composition across age–type groups, and

⁵¹Aggregate productivity increases by the same percentage as consumption because the labor endowment is fixed. Stationary firm mass reflects both entry and exit.

within-group performance. The first two capture adjustments in the firm population, while within-group performance measures mean deflated sales per firm and incorporates both production and portfolio adjustments.⁵² Firm composition and within-group performance contribute 6.26 and 2.29 percentage points, respectively, offsetting the 8.01 percentage-point reduction associated with lower firm mass.

To separate composition and performance from changes in firm mass, we normalize both distributions to baseline firm mass. Under this accounting normalization, aggregate deflated sales rise by 8.90%. Firm composition contributes 6.51 percentage points (73.21%), while within-group performance contributes 2.38 percentage points (26.79%). Online Appendix [OA.C.4](#) details the decomposition and resource accounts. Improved information processing thus raises aggregate consumption through changes in firm composition and performance even though fewer firms operate.

4.8 Welfare Under a Restricted Social Planner

We next consider a restricted planner choosing pre-entry capability investment, production, product switching, startup operation, and entry to maximize welfare. The planner respects physical feasibility and retains the CE's value-function restrictions on product switching, startup operation, and entry. Assigned endogenous-exit values create no resources; welfare depends on household consumption. Appendix [C](#) formalizes the benchmark.

For fixed capabilities and acquisition costs, the restricted planner's admissible stationary allocations coincide with stationary CE allocations; consumption and welfare coincide under a common equilibrium selection (Proposition [C.1](#)). Optimizing capability investment weakly raises attainable welfare because the baseline CE remains feasible (Corollary [C.1](#)).

4.9 Discussion and Policy Implications

Our framework links the up-or-out dynamics of young firms to learning about what to produce. Firms learn about expertise in individual products, so growth reflects both changes in

⁵²Within-group performance also reflects equilibrium price adjustments and selection within each group. The accounting components therefore do not map one for one to the margins in the paired experiment.

production and changes in the products they operate, while endogenous exit occurs when no product is retained. Information-processing ability governs how precisely firms make these decisions. Managerial investment before entry therefore affects later growth and exit through information, without directly changing productivity or product-specific expertise.

Intercohort capability spillovers provide a motive for intervention because entrants do not internalize the benefits of their investment to future cohorts. Evidence that managerial and entrepreneurial capabilities can be learned and transmitted supports this channel (Giorcelli, 2019; Guiso et al., 2021). The appropriate mix of managerial and productivity investment depends on their returns, acquisition costs, and transferability.

Management training, mentoring, and access to data are potential ways to improve information processing. Our subsidy experiment shows why its benefits need not appear in survival rates: improved information raises cumulative sales through production and portfolio adjustments even as more entrants exit, particularly those with weak initial products. In general equilibrium, consumption rises after accounting for the subsidy's resource cost. Evaluating capability investment therefore requires considering how firms grow and select products, as well as how many survive.

5 Concluding Remarks

Young firms must learn what to produce as well as how much to produce. We develop a model in which firms learn about their product-specific expertise and differ in how precisely they process information. Product selection connects this learning to firm growth and exit. In U.S. manufacturing, selection on expertise is weaker among young firms and stronger among firms with higher management scores, greater data availability, and more intensive data use, consistent with the model's mechanism.

The calibrated model shows how better information can strengthen both growth and exit. A subsidy to startups' managerial investment increases cumulative sales through production and portfolio adjustments while reducing survival, with larger early exit responses among firms that begin with weak products. In general equilibrium, consumption rises by 0.54% net of the subsidy's resource cost even as fewer firms operate. Survival alone therefore

understates the benefits of improved information processing.

With capabilities fixed and a common equilibrium selection, the restricted planner replicates the stationary competitive equilibrium. Capability spillovers across cohorts provide a rationale for intervention. Policies toward young firms should account for how well firms learn what to produce, alongside their effects on entry, growth, and survival.

References

- Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr**, “Innovation, reallocation, and growth,” *American Economic Review*, 2018, 108 (11), 3450–3491.
- Akcigit, Ufuk and Sina T Ates**, “What Happened to US Business Dynamism?,” *Journal of Political Economy*, 2023, 131 (8), 2059–2124.
- **and William R. Kerr**, “Growth through Heterogeneous Innovations,” *Journal of Political Economy*, 2018, 126 (4), 1374–1443.
- , **Harun Alp, and Michael Peters**, “Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries,” *American Economic Review*, 2021, 111 (1), 231–275.
- Alon, Titan, David Berger, Robert Dent, and Benjamin Pugsley**, “Older and slower: The startup deficit’s lasting effects on aggregate productivity growth,” *Journal of Monetary Economics*, 2018, 93, 68–85.
- Argente, David, Doireann Fitzgerald, Sara Moreira, and Anthony Priolo**, “How Do Entrants Build Market Share? The Role of Demand Frictions,” *American Economic Review: Insights*, 2025, 7 (2), 160–176.
- , **Munseob Lee, and Sara Moreira**, “Innovation and product reallocation in the great recession,” *Journal of Monetary Economics*, 2018, 93, 1–20.
- , — , **and —** , “The life cycle of products: Evidence and implications,” *Journal of Political Economy*, 2024, 132 (2), 337–390.
- , **Sara Moreira, Ezra Oberfield, and Venky Venkateswaran**, “Scalable Expertise: How Standardization Drives Scale and Scope,” Working Paper 34160, National Bureau of Economic Research August 2025.
- Arkolakis, Costas, Theodore Papageorgiou, and Olga A Timoshenko**, “Firm learning and growth,” *Review of Economic Dynamics*, 2018, 27, 146–168.
- Asriyan, Vladimir and Alexandre N. Kohlhas**, “The Macroeconomics of Data: Scale, Product Choice, and Pricing in the Information Age,” 2026. Working paper, March.
- Ates, Sina T. and Felipe E. Saffie**, “Fewer but Better: Sudden Stops, Firm Entry, and Financial Selection,” *American Economic Journal: Macroeconomics*, 2021, 13 (3), 304–356.
- Baley, Isaac, Ana Figueiredo, and Robert Ulbricht**, “Mismatch Cycles,” *Journal of Political Economy*, 2022, 130 (11), 2943–2984.
- **and Laura L Veldkamp**, *The Data Economy: Tools and Applications*, Princeton University Press, 2025.
- Becker, Randy, Wayne Gray, and Jordan Marvakov**, “NBER-CES manufacturing industry database: Technical notes,” *NBER Working Paper*, 2013, 5809.
- Begenau, Juliane, Maryam Farboodi, and Laura Veldkamp**, “Big data in finance and the growth of large firms,” *Journal of Monetary Economics*, 2018, 97, 71–87.
- Berlingieri, Giuseppe, Maarten De Ridder, Danial Lashkari, and Davide Rigo**, “Creative Destruction through Innovation Bursts,” Discussion Paper 21795, Centre for Economic Policy Research July 2026.
- Berman, Nicolas, Vincent Rebeyrol, and Vincent Vicard**, “Demand learning and firm dynamics: evidence from exporters,” *Review of Economics and Statistics*, 2019, 101 (1), 91–106.
- Bernard, Andrew B, Stephen J Redding, and Peter K Schott**, “Multiple-product firms and

- product switching,” *American Economic Review*, 2010, 100 (1), 70–97.
- , —, and —, “Multiproduct firms and trade liberalization,” *Quarterly Journal of Economics*, 2011, 126 (3), 1271–1318.
- Bloom, Nicholas, Benn Eifert, Aprajit Mahajan, David McKenzie, and John Roberts**, “Does Management Matter? Evidence from India,” *Quarterly Journal of Economics*, 2013, 128 (1), 1–51.
- , **Erik Brynjolfsson, Lucia Foster, Ron Jarmin, Megha Patnaik, Itay Saporta-Eksten, and John Van Reenen**, “What drives differences in management practices?,” *American Economic Review*, 2019, 109 (5), 1648–1683.
- , **Takafumi Kawakubo, Charlotte Meng, Paul Mizen, Rebecca Riley, Tatsuro Senga, and John Van Reenen**, “Do Well Managed Firms Make Better Forecasts?,” *Review of Economics and Statistics*, 2025, pp. 1–26.
- Boehm, Johannes, Swati Dhingra, and John Morrow**, “The comparative advantage of firms,” *Journal of Political Economy*, 2022, 130 (12), 3025–3100.
- Bonhomme, Stéphane, Thibaut Lamadon, and Elena Manresa**, “A distributional framework for matched employer employee data,” *Econometrica*, 2019, 87 (3), 699–739.
- Broda, Christian and David E Weinstein**, “Globalization and the Gains from Variety,” *Quarterly Journal of Economics*, 2006, 121 (2), 541–585.
- and —, “Product Creation and Destruction: Evidence and Price Implications,” *American Economic Review*, June 2010, 100 (3), 691–723.
- Bruhn, Miriam, Dean Karlan, and Antoinette Schoar**, “The Impact of Consulting Services on Small and Medium Enterprises: Evidence from a Randomized Trial in Mexico,” *Journal of Political Economy*, 2018, 126 (2), 635–687.
- Buera, Francisco J. and Ezra Oberfield**, “The Global Diffusion of Ideas,” *Econometrica*, 2020, 88 (1), 83–114.
- Camuffo, Arnaldo, Alessandro Cordova, Alfonso Gambardella, and Chiara Spina**, “A Scientific Approach to Entrepreneurial Decision Making: Evidence from a Randomized Control Trial,” *Management Science*, 2020, 66 (2), 564–586.
- Clementi, Gian Luca and Hugo A Hopenhayn**, “A theory of financing constraints and firm dynamics,” *Quarterly Journal of Economics*, 2006, 121 (1), 229–265.
- Cooley, Thomas F and Vincenzo Quadrini**, “Financial markets and firm dynamics,” *American Economic Review*, 2001, 91 (5), 1286–1310.
- David, Joel M., Hugo A. Hopenhayn, and Venky Venkateswaran**, “Information, Misallocation, and Aggregate Productivity,” *Quarterly Journal of Economics*, 2016, 131 (2), 943–1005.
- Decker, Ryan, John Haltiwanger, Ron Jarmin, and Javier Miranda**, “The role of entrepreneurship in US job creation and economic dynamism,” *Journal of Economic Perspectives*, 2014, 28 (3), 3–24.
- Dent, Robert C, Benjamin W Pugsley, Harrison Wheeler et al.**, “Longitudinal Linking of Enterprises in the LBD and SSL,” Technical Report, Center for Economic Studies, US Census Bureau 2018.
- Dunne, Timothy, Mark J. Roberts, and Larry Samuelson**, “Patterns of Firm Entry and Exit in U.S. Manufacturing Industries,” *RAND Journal of Economics*, 1988, 19 (4), 495–515.
- , —, and —, “The Growth and Failure of U.S. Manufacturing Plants,” *Quarterly Journal of Economics*, 1989, 104 (4), 671–698.

- , **Shawn D. Klimek**, and **Mark J. Roberts**, “Exit from Regional Manufacturing Markets: The Role of Entrant Experience,” *International Journal of Industrial Organization*, 2005, 23 (5–6), 399–421.
- Eckel, Carsten** and **J Peter Neary**, “Multi-product firms and flexible manufacturing in the global economy,” *Review of Economic Studies*, 2010, 77 (1), 188–217.
- Eslava, Marcela**, **John C. Haltiwanger**, and **Nicolás Urdaneta**, “The Size and Life-Cycle Growth of Plants: The Role of Productivity, Demand, and Wedges,” *Review of Economic Studies*, 2024, 91 (1), 259–300.
- Farboodi, Maryam** and **Laura Veldkamp**, “A Model of the Data Economy,” *Review of Economic Studies*, 2026. Advance online publication.
- , **Roxana Mihet**, **Thomas Philippon**, and **Laura Veldkamp**, “Big Data and Firm Dynamics,” *AEA Papers and Proceedings*, 2019, 109, 38–42.
- Foster, Lucia**, **John Haltiwanger**, and **Chad Syverson**, “Reallocation, firm turnover, and efficiency: selection on productivity or profitability?,” *American Economic Review*, 2008, 98 (1), 394–425.
- , —, and —, “The slow growth of new plants: learning about demand?,” *Economica*, 2016, 83 (329), 91–129.
- Giorcelli, Michela**, “The Long-Term Effects of Management and Technology Transfers,” *American Economic Review*, 2019, 109 (1), 121–152.
- Gourio, François** and **Leena Rudanko**, “Customer Capital,” *Review of Economic Studies*, 2014, 81 (3), 1102–1136.
- Groes, Fane**, **Philipp Kircher**, and **Iourii Manovskii**, “The U-Shapes of Occupational Mobility,” *Review of Economic Studies*, 2015, 82 (2), 659–692.
- Guiso, Luigi**, **Luigi Pistaferri**, and **Fabiano Schivardi**, “Learning Entrepreneurship from Other Entrepreneurs?,” *Journal of Labor Economics*, 2021, 39 (1), 135–191.
- Guzman, Jorge** and **Scott Stern**, “The State of American Entrepreneurship: New Estimates of the Quantity and Quality of Entrepreneurship for 32 US States, 1988–2014,” *American Economic Journal: Economic Policy*, 2020, 12 (4), 212–243.
- Haltiwanger, John**, **Ron S Jarmin**, and **Javier Miranda**, “Who creates jobs? Small versus large versus young,” *Review of Economics and Statistics*, 2013, 95 (2), 347–361.
- , **Ron S. Jarmin**, **Robert B. Kulick**, and **Javier Miranda**, “High-Growth Young Firms: Contribution to Job, Output, and Productivity Growth,” in John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar, eds., *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, Studies in Income and Wealth, Chicago: University of Chicago Press, 2017, pp. 11–62.
- Hausmann, Ricardo** and **Dani Rodrik**, “Economic Development as Self-Discovery,” *Journal of Development Economics*, 2003, 72 (2), 603–633.
- Hopenhayn, Hugo A.**, “Entry, Exit, and Firm Dynamics in Long Run Equilibrium,” *Econometrica*, 1992, 60 (5), 1127–1150.
- Hottman, Colin J.**, **Stephen J. Redding**, and **David E. Weinstein**, “Quantifying the Sources of Firm Heterogeneity,” *Quarterly Journal of Economics*, 2016, 131 (3), 1291–1364.
- Hsieh, Chang-Tai** and **Peter J. Klenow**, “The Life Cycle of Plants in India and Mexico,” *Quarterly Journal of Economics*, 2014, 129 (3), 1035–1084.
- Hurst, Erik** and **Benjamin Wild Pugsley**, “What Do Small Businesses Do?,” *Brookings*

- Papers on Economic Activity*, 2011, 42 (2), 73–142.
- Jaef, Roberto N. Fattal**, “Entry and Exit, Multiproduct Firms, and Allocative Distortions,” *American Economic Journal: Macroeconomics*, 2018, 10 (2), 86–112.
- Jo, Karam and Seula Kim**, “Heterogeneous innovations and growth under imperfect technology spillovers,” *Journal of Monetary Economics*, 2026, 158, 103899.
- Jones, Charles I.**, “R&D-Based Models of Economic Growth,” *Journal of Political Economy*, 1995, 103 (4), 759–784.
- Jovanovic, Boyan**, “Job Matching and the Theory of Turnover,” *Journal of Political Economy*, 1979, 87 (5, Part 1), 972–990.
- , “Selection and the Evolution of Industry,” *Econometrica*, 1982, 50 (3), 649–670.
- Kehrig, Matthias and Nicolas Vincent**, “Good Dispersion, Bad Dispersion,” *Review of Economic Studies*, 2026, 93 (2), 1103–1136.
- Kerr, William R, Ramana Nanda, and Matthew Rhodes-Kropf**, “Entrepreneurship as experimentation,” *Journal of Economic Perspectives*, 2014, 28 (3), 25–48.
- Kim, Seula**, “Workers’ Job Prospects and Young Firm Dynamics,” 2026. Working paper, Pennsylvania State University.
- Klette, Tor Jakob and Samuel Kortum**, “Innovating firms and aggregate innovation,” *Journal of Political Economy*, 2004, 112 (5), 986–1018.
- Koning, Rembrand, Sharique Hasan, and Aaron Chatterji**, “Experimentation and Start-up Performance: Evidence from A/B Testing,” *Management Science*, 2022, 68 (9), 6434–6453.
- Lee, Yi, Shengyu Li, and Mark J. Roberts**, “Factor Bias and Heterogeneous Output Quality in Multiproduct Production,” Working Paper 35274, National Bureau of Economic Research May 2026.
- Luttmer, Erzo G. J.**, “Selection, Growth, and the Size Distribution of Firms,” *Quarterly Journal of Economics*, 2007, 122 (3), 1103–1144.
- Mayer, Thierry, Marc J Melitz, and Gianmarco IP Ottaviano**, “Market size, competition, and the product mix of exporters,” *American Economic Review*, 2014, 104 (2), 495–536.
- Mitchell, Matthew F**, “The scope and organization of production: firm dynamics over the learning curve,” *RAND Journal of Economics*, 2000, 31 (1), 180–205.
- Papageorgiou, Theodore**, “Learning Your Comparative Advantages,” *Review of Economic Studies*, 2014, 81 (3), 1263–1295.
- Perla, Jesse and Christopher Tonetti**, “Equilibrium Imitation and Growth,” *Journal of Political Economy*, 2014, 122 (1), 52–76.
- Queiró, Francisco**, “Entrepreneurial Human Capital and Firm Dynamics,” *Review of Economic Studies*, 2022, 89 (4), 2061–2100.
- Romer, Paul M.**, “Endogenous Technological Change,” *Journal of Political Economy*, 1990, 98 (5, Part 2), S71–S102.
- Sampson, Thomas**, “Dynamic Selection: An Idea Flows Theory of Entry, Trade, and Growth,” *Quarterly Journal of Economics*, 2016, 131 (1), 315–380.
- Schmalz, Martin C, David A Sraer, and David Thesmar**, “Housing collateral and entrepreneurship,” *The Journal of Finance*, 2017, 72 (1), 99–132.
- Sedláček, Petr and Vincent Sterk**, “The growth potential of startups over the business cycle,” *American Economic Review*, 2017, 107 (10), 3182–3210.
- Sterk, Vincent, Petr Sedláček, and Benjamin Pugsley**, “The nature of firm growth,” *Amer-*

- ican Economic Review*, 2021, 111 (2), 547–579.
- Stokey, Nancy L.**, *The Economics of Inaction: Stochastic Control Models with Fixed Costs*, Princeton, NJ: Princeton University Press, 2009.
- Stokey, Nancy L and Robert E Lucas Jr**, *Recursive methods in economic dynamics*, Harvard University Press, 1989.
- Timoshenko, Olga A**, “Product switching in a model of learning,” *Journal of International Economics*, 2015, 95 (2), 233–249.

A Theory: Proof Sketches and Supporting Derivations

This appendix sketches the proofs of the theoretical results in Section 2, in the order in which they appear. Unless stated otherwise, aggregate prices and quantities are stationary, $0 < \beta < 1$, $0 \leq \delta < 1$, $0 < x < 1$, $\eta > 1$, and prior and signal variances are positive and finite. Comparisons condition on firm types and the specified information-processing path, and endogenous-exit payoffs follow the convention in Section 2.3.7.

A.1 Learning and Bayesian Shrinkage

Proof sketch of Lemma 1. Fix true expertise and the observation window, without conditioning on product retention. Taking expectations in (2.3) and using (2.4) gives

$$\mathbb{E}[\hat{\theta}_{ipt} \mid \theta_{ip}] - \log \theta_{ip} = \frac{\hat{\sigma}_{ipt}^2}{\sigma_0^2} (\theta_0 - \log \theta_{ip}), \quad \frac{\partial \hat{\sigma}_{ipt}^2}{\partial \psi_{it}} = -\frac{(\hat{\sigma}_{ipt}^2)^2}{\sigma_\varepsilon^2} < 0.$$

The bias formula and variance derivative establish the result. Since processing enters posterior variance through its cumulative sum, the comparison also holds for lower cumulative processing. Online Appendix OA.A.1 gives the full proof. $\square$

A.2 Existence, Uniqueness, and Ordering of Thresholds

Proof sketch of Proposition 1. Values and monotonicity. The posterior moment $H_{it}(\hat{\theta}, \tau) = \exp\{(\eta-1)\hat{\theta} + (\eta-1)^2 \hat{\sigma}_{ipt}^2 / 2\}$ is a martingale under retention and resets to its finite prior value H^0 upon replacement. Expected absolute payoffs j periods ahead are bounded, uniformly over feasible policies, by a constant multiple of $1 + H_{it} + (j+1)H^0$. Discounting therefore gives a vanishing payoff tail and compact-uniform convergence of finite-horizon values. Coupling posterior transitions with the same normal innovation shows the Bellman operator preserves monotonicity. Continuous and strictly increasing current profits make V_{it}^{ns} and V_{it}^{cs} continuous and strictly increasing in $\hat{\theta}$, while V_{it} is continuous and weakly increasing.

Crossings and ordering. Direct subtraction of the action-value equations gives

$$V_{it}^{cs}(\hat{\theta}, \tau) - V_{it}^{ds}(\emptyset) = (1-x) \left[V_{it}^{ns}(\hat{\theta}, \tau) - \beta(1-\delta)V_{i,t+1}^{ds}(\emptyset) \right],$$

$$V_{it}^{ns}(\hat{\theta}, \tau) - V_{it}^{cs}(\hat{\theta}, \tau) = w_t f^s + x \left[V_{it}^{ns}(\hat{\theta}, \tau) - V_{it}^{ns}(\hat{\theta}^0, 0) \right]. \quad (\text{A.1})$$

Under (2.18), the occupied line's surplus over vacancy is bounded by a constant multiple of H_{it} . Hence

$$\lim_{\hat{\theta} \rightarrow -\infty} V_{it}^{ns}(\hat{\theta}, \tau) = -\beta w_t f^o + \beta(1 - \delta)V_{i,t+1}^{ds}(\emptyset).$$

The lower value-matching level is $\beta(1 - \delta)V_{i,t+1}^{ds}(\emptyset)$. The vacancy equation implies that the upper level, $V_{it}^{ns}(\hat{\theta}^0, 0) - w_t f^s/x$, exceeds it by $[V_{it}^{ds}(\emptyset) - \beta(1 - \delta)V_{i,t+1}^{ds}(\emptyset)]/x > 0$. Thus both levels exceed the low-belief limit. Since V_{it}^{ns} tends to infinity at high beliefs, continuity and strict monotonicity give unique ordered crossings. Equation (A.1) establishes the policy regions, and monotonicity of $\mathcal{B}_{it}$ gives the demand cutoffs. Online Appendix OA.A.3.2 supplies the full argument and a primitive sufficient condition. $\square$

The drop condition also has a useful economic interpretation:

$$\bar{\pi}_{it+1}(\hat{\theta}_{it}^{ds}(\tau), \tau) = -(1 - \delta)\mathbb{E}_{it} \left[V_{i,t+1}(\hat{\theta}', \tau + 1) - V_{i,t+1}^{ds}(\emptyset) \right] \leq 0, \quad (\text{A.2})$$

where the expectation is evaluated at the drop cutoff. Retaining the incumbent preserves its continuation surplus over a vacancy, allowing the firm to tolerate negative expected operating profit.

A.3 Destination Assignment and Product-Line Separability

Destination assignment. The restriction $n_i \leq \lfloor N/2 \rfloor$ makes every subset of searching lines feasible. Uniform sampling without replacement gives each line marginal destination probability $1/(N - n_i)$; symmetric markets and the common-prior reset make all feasible destinations equally valuable. At a startup's first search, only its initial product is excluded. Online Appendix OA.A.2 gives the construction and proof.

Proof sketch of Lemma 2 and Proposition 2. From age two on, substitute the sum of product-line values into the firm's recursion. If at least one product is retained, additivity and linearity of expectation give the sum of the chosen line-level action values. If all products are dropped, payoff equivalence gives $n_i V_{it}^{ds}(\emptyset)$, the sum of line-level drop values. Maximizing

each line's value therefore maximizes firm value. The single-line payoff bound verifies this recursion over the infinite horizon. Online Appendix [OA.A.4](#) gives the full proof. $\square$

A.4 Distributions and Aggregate Accounting

Product-line measures. Equations (2.30) and (2.31) give the exact firm transition and active-line projection. For computation, we include vacancies and approximate the joint distribution of belief, tenure, and vacancy states by the product of their common marginals conditional on firm age and types. With $q^\emptyset$ the post-drop vacancy probability, a dropped or vacant line reaches search with probability $1 - p^e$, where $p^e = (q^\emptyset)^{n_i-1}$ for $n_i > 1$ and $p^e = 1$ for $n_i = 1$. At age one, continuation depends exactly on retaining the startup's initial product.

Conditional on exogenous survival, retained products remain in their markets after no search or failed replacement search. Successful search resets beliefs to the common prior; failed search leaves vacant lines vacant. Symmetry gives destination probability $1/(N - 1)$ after averaging over other excluded products and assigns entrants uniformly across markets. Operating costs count producing lines before the death shock, search costs count searches after both exit decisions, and entry costs count all entry attempts. Online Appendix [OA.A.5](#) gives the transition equations.

Price index. Substituting (2.9) into the CES price index and using $\mathbb{E}_{t-1}[\theta_{ipt}^{(\eta-1)/\eta}] = \mathcal{B}(\xi)$ gives

$$P_{pt}^{1-\eta} = \left(\frac{\eta}{\eta-1} w_t \right)^{1-\eta} \int \mathcal{B}(\xi)^\eta z_i^{\eta-1} \mu_{p,t-1}(d\xi). \quad (\text{A.3})$$

This proves (2.36). Combining it with variable labor demand and labor-market clearing gives (2.37). Online Appendix [OA.A.6](#) gives the search-cost expression and full derivation.

A.5 Output Wedges

Proof sketch of Proposition 3. With known permanent expertise, $\mathcal{B}^{FI} = \theta_{ip}^{(\eta-1)/\eta} \exp\{(\eta - 1)^2 \sigma_\varepsilon^2 / (2\eta^2)\}$. The optimal quantity ratio is $(\mathcal{B}/\mathcal{B}^{FI})^\eta$, giving (2.38). Conditional on true expertise, the posterior mean is normal with the mean in Appendix [A.1](#) and variance $\hat{\sigma}^2(1 -$

$\hat{\sigma}^2/\sigma_0^2$), where $\hat{\sigma}^2 \equiv \hat{\sigma}_{ip,t-1}^2$. Its exponential moment gives

$$\log \mathbb{E}[1 - \omega_{ipt} \mid \theta_{ip}] = \frac{(\eta - 1)\hat{\sigma}^2}{\sigma_0^2} \left\{ \theta_0 - \log \theta_{ip} + \frac{\eta - 1}{2} \left[\left(1 + \frac{1}{\eta}\right) \sigma_0^2 - \hat{\sigma}^2 \right] \right\}.$$

Expected underproduction occurs when the term in braces is negative. The cutoff exceeds the prior mean because posterior uncertainty and the dispersion of posterior means both raise expected output through its nonlinearity. Higher posterior variance lowers this term and makes the whole expression more negative wherever underproduction occurs, proving both comparisons. Online Appendix [OA.A.7](#) gives the full derivation. $\square$

A.6 Convexity and the Value of Further Learning

Lemma A.1 (Value of Further Learning). *At fixed tenure and firm states, the product-line value is continuous, weakly increasing, and convex in the demand term $\mathcal{B}$, and the value of retaining without search is strictly increasing and strictly convex. Holding current $\mathcal{B}$ and future values fixed, a larger variance of the next belief update raises expected continuation value.*

Proof sketch. Let $\alpha = (\eta - 1)/\eta$ and let $v_t^\tau = \hat{\sigma}_{ipt}^2 - \hat{\sigma}_{ip,t+1}^2$ be the variance of the next belief update. Gaussian updating implies the retained-product transition

$$\mathcal{B}' = \mathcal{B} \exp\{\alpha \sqrt{v_t^\tau} u - \alpha^2 v_t^\tau / 2\}, \quad u \sim N(0, 1), \quad \mathbb{E}_t \mathcal{B}' = \mathcal{B}. \quad (\text{A.4})$$

Expected profit is strictly increasing and strictly convex in $\mathcal{B} > 0$. Since $\mathcal{B}'$ is a positive multiple of $\mathcal{B}$, taking expectations preserves monotonicity and convexity. So does the maximum over actions, because the continue-and-search value is $(1 - x)V^{ns}$ plus a term independent of $\mathcal{B}$, and the drop-and-search value does not depend on $\mathcal{B}$. Backward induction and the payoff bound in Appendix [A.2](#) give the stated properties.

A larger v_t^τ is equivalent to multiplying $\mathcal{B}'$ by an independent lognormal factor with mean one. By conditional Jensen's inequality, this raises the expectation of any convex continuation value, strictly when the value is strictly convex on a set reached with positive probability. Online Appendix [OA.A.8](#) gives the full proof. $\square$

A.7 Information Processing, Productivity, and Product Tenure

We use the separable product-line problem and the maintained assumptions of Proposition 1, at firm ages two and above. Normalize the stationary wage to one and write $d = \beta(1 - \delta)$. Comparisons hold aggregate prices, product scope, search technology, and fixed costs fixed. The proofs use value matching and backward induction, following Stokey and Lucas Jr (1989, Sections 9.1 and 9.4), and the convexity argument in Stokey (2009, Theorem 4.1). Online Appendix OA.A.9 supplies the full proofs.

The two decision margins. Write $D_t = V_{it}^{ds}(\emptyset)$ and define

$$\begin{aligned} U_t(\mathcal{B}, \tau) &= V_{it}^{ns}(\mathcal{B}, \tau) - dD_{t+1}, \\ \Delta_t^V &= V_{it}^{ns}(\mathcal{B}^0, 0) - dD_{t+1} - f^s/x = (D_t - dD_{t+1})/x. \end{aligned} \quad (\text{A.5})$$

Here U_t compares retaining the existing product without search today with a vacant line from next period, and Δ_t^V is the corresponding value of a new product net of search cost per successful replacement. The thresholds satisfy

$$U_t(\mathcal{B}_{it}^{ds}, \tau) = 0, \quad U_t(\mathcal{B}_{it}^{cs}, \tau) = \Delta_t^V. \quad (\text{A.6})$$

The first equality follows because both dropping and retaining while searching involve the same current search opportunity; their difference is the value of keeping the existing product if search fails. The second equates the gain from a successful replacement to its search cost. The Bellman equations give

$$\begin{aligned} U_t(\mathcal{B}, \tau) &= \beta \bar{\pi}_{it+1}(\mathcal{B}, z_i) \\ &\quad + d\mathbb{E}_t \max\{U_{t+1}(\mathcal{B}', \tau + 1) - x\Delta_{t+1}^V, (1 - x)U_{t+1}(\mathcal{B}', \tau + 1), 0\}. \end{aligned} \quad (\text{A.7})$$

The maximum represents retaining without search, retaining while searching, and dropping and searching. It is increasing in the value of the existing product and convex in its demand term. Also, $U_{t,\mathcal{B}} \geq \beta \bar{\pi}_{it+1,\mathcal{B}} > 0$. Thus, a lower U_t raises both thresholds when Δ_t^V

is unchanged.

The productivity conditions. Let $\mathcal{K}_{it}^{ns}(\mathcal{B}, \tau)$ be the expected discounted sum of operating and search fixed costs when the firm retains the product without search today and makes optimal decisions thereafter. Let $\mathcal{K}_{it}^{ds}(\emptyset)$ be the corresponding sum when the firm drops the product and searches today. Both cost burdens include operating and search fixed costs over all subsequent periods and product replacements.⁵³ The superscripts identify today's decision, not the types of costs included or the firm's future decisions. In particular, $\mathcal{K}^{ns}$ begins with βf^o and includes future search costs; $\mathcal{K}^{ds}$ begins with f^s and includes operating costs when search succeeds. The conditions for part (ii) are

$$\mathcal{K}_{it}^{ns}(\mathcal{B}_{it}^{ds}, \tau) > d\mathcal{K}_{i,t+1}^{ds}(\emptyset), \quad (\text{A.8})$$

$$\mathcal{K}_{it}^{ns}(\mathcal{B}_{it}^{cs}, \tau) - \mathcal{K}_{it}^{ns}(\mathcal{B}^0, 0) < f^s/x. \quad (\text{A.9})$$

Each applies at its corresponding threshold. For a finite productivity comparison, impose the respective condition throughout the comparison interval. These are conditions on expected costs under optimal decisions; the Online Appendix also gives sufficient restrictions stated entirely in primitives.

Proof sketch of Proposition 4. Information processing. Hold processing at all dates other than t fixed and take $\tau \geq 1$. The posterior variance of the existing product is

$$\hat{\sigma}_{ipt}^2 = \left[\sigma_0^{-2} + \frac{\tau + \sum_{s=t-\tau+1}^t \psi_{is}}{\sigma_\varepsilon^2} \right]^{-1}. \quad (\text{A.10})$$

Higher current processing raises the precision of this posterior and of every subsequent posterior while the product is retained. With future signals unchanged, the variance of each subsequent belief revision falls:

$$\frac{\partial v_{t+j}^{\tau+j}}{\partial \psi_{it}} = -\frac{\hat{\sigma}_{ip,t+j}^4 - \hat{\sigma}_{ip,t+j+1}^4}{\sigma_\varepsilon^2} < 0, \quad j \geq 0, \quad (\text{A.11})$$

⁵³These sums follow the same accounting as the value functions, including the fixed-cost component of the assigned vacant-line value at endogenous exit. They exclude variable production costs and pre-entry investment costs.

where $v_{t+j}^{\tau+j} = \hat{\sigma}_{ip,t+j}^2 - \hat{\sigma}_{ip,t+j+1}^2$ and $\hat{\sigma}^4 = (\hat{\sigma}^2)^2$. New products begin at the common prior and receive their first observations next period, so their values and Δ_{t+j}^V are unchanged. At a given current demand term, smaller belief revisions reduce the value of future production and switching decisions by the learning-value lemma. Applying this argument backward through (A.7) gives $U_{t,\psi_{it}} < 0$. The strict inequality follows because further learning affects production on histories reached with positive probability. Differentiating (A.6) therefore gives $\partial \mathcal{B}_{it}^j / \partial \psi_{it} = -U_{t,\psi_{it}} / U_{t,B} > 0$ for $j \in \{ds, cs\}$, evaluated at the corresponding threshold. Discounting makes the expected contribution of sufficiently distant payoffs arbitrarily small, so the backward argument applies to the infinite-horizon problem.

Productivity. Hold the entire processing path fixed. For a given product-switching policy, with production quantities chosen optimally, gross operating payoffs—revenues less variable production costs—are proportional to $z_i^{\eta-1}$. Fixed costs do not change with productivity. The envelope theorem therefore gives

$$\frac{z_i}{\eta - 1} V_{it,z_i}^j = V_{it}^j + \mathcal{K}_{it}^j, \quad j \in \{ns, ds\}. \quad (\text{A.12})$$

The right-hand side is the expected discounted sum of gross operating payoffs under optimal decisions, using the same accounting for the assigned continuation value at endogenous exit. At the drop threshold, $V_{it}^{ns}(\mathcal{B}_{it}^{ds}, \tau) = dD_{t+1}$. Condition (A.8) therefore says that retaining the existing product generates greater gross operating payoffs than a vacant line from next period: its larger fixed costs must be offset to give the same net value. Higher productivity favors retention, lowering the drop threshold. At the search threshold, $V_{it}^{ns}(\mathcal{B}_{it}^{cs}, \tau) = V_{it}^{ns}(\mathcal{B}^0, 0) - f^s/x$. Condition (A.9) says that a new product generates greater gross operating payoffs than the existing product. Higher productivity increases this advantage, making search worthwhile at a higher demand term. Differentiating the two threshold equations using (A.12) gives the stated signs because $V_{it,B}^{ns} > 0$ at both thresholds.

Product tenure. At the same firm age and processing path, a product of tenure $\tau + 1$ has the additional observation dated $t - \tau$. This raises inherited precision by $(1 + \psi_{i,t-\tau}) / \sigma_\varepsilon^2 > 0$. Less remains to be learned about the existing product, while future processing and the value of new products are unchanged. The same argument as in part (i) lowers U_t at each current

demand term and raises both thresholds. □

The gap between thresholds. Under both productivity conditions, the search threshold rises and the drop threshold falls, so $\partial[\mathcal{B}_{it}^{cs}(\tau) - \mathcal{B}_{it}^{ds}(\tau)]/\partial z_i > 0$. The processing and tenure results raise both thresholds but do not determine which moves farther.

A.8 Returns to Pre-Entry Investment

Hold aggregate prices, quantities, and the joint distribution of normalized types $(\varphi_i, \tilde{m}_i)$ fixed, with $z_i = z_0 \varphi_i$ and $m_i = m_0 \tilde{m}_i$. Write expected entry value before investment costs as $\mathcal{E} = \mathcal{G} - \mathcal{K}$, where $\mathcal{G}$ is expected discounted revenue less variable production costs and $\mathcal{K}$ is expected discounted operating and search fixed costs, valued at fixed wages. Both components include replacements and assigned vacant-line payoffs at endogenous exit.

Higher m_0 permits replication of lower-precision signals by adding independent noise, giving $\mathcal{E}_{m_0} \geq 0$. For a fixed operating and switching policy, including assigned exit continuations, optimized gross operating payoffs scale with $z_0^{\eta-1}$. The envelope theorem gives, where derivatives exist and differentiation under expectations is valid,

$$\mathcal{E}_{z_0} = \frac{\eta-1}{z_0} \mathcal{G}, \quad \mathcal{E}_{z_0 m_0} = \frac{\eta-1}{z_0} \mathcal{G}_{m_0} = \frac{\eta-1}{z_0} (\mathcal{E}_{m_0} + \mathcal{K}_{m_0}). \quad (\text{A.13})$$

The productivity derivative is positive if production occurs with positive probability, and local complementarity holds if and only if $\mathcal{G}_{m_0} \geq 0$. These are total derivatives under optimal policies: at indifference, boundary terms cancel in net value but need not cancel in either component separately. Online Appendix [OA.A.13](#) gives the full proof.

B Persistent Demand and the Expertise Measure

B.1 Persistent Product-Market Shocks

Quasi-differencing gives $\log \theta_{ip,t+1} - \rho \log \theta_{ipt} = (1 - \rho) \log \theta_{ip} + e_{ip,t+1}$, where $e_{ip,t+1}$ has variance $(1 - \rho^2) \sigma_\varepsilon^2$ and is independent of past information and private-signal errors. For

positive tenure, normal conjugacy gives

$$\begin{aligned}\hat{\sigma}_{ip,t+1}^{-2} &= \hat{\sigma}_{ipt}^{-2} + \frac{1-\rho}{(1+\rho)\sigma_\varepsilon^2} + \frac{\psi_{i,t+1}}{\sigma_\varepsilon^2}, \\ \hat{\theta}_{ip,t+1} &= \hat{\sigma}_{ip,t+1}^2 \left[\frac{\hat{\theta}_{ipt}}{\hat{\sigma}_{ipt}^2} + \frac{\log \theta_{ip,t+1} - \rho \log \theta_{ipt}}{(1+\rho)\sigma_\varepsilon^2} + \frac{\psi_{i,t+1} \log s_{ip,t+1}}{\sigma_\varepsilon^2} \right].\end{aligned}\quad (\text{B.1})$$

A product's first market observation has precision $1/\sigma_\varepsilon^2$, because a new product starts with a stationary demand shock, so its first update is the baseline update from (θ_0, σ_0^2) .

For positive tenure, conditional on current information, $\log \theta_{ip,t+1}$ is normal with mean $\rho \log \theta_{ipt} + (1-\rho)\hat{\theta}_{ipt}$ and variance $(1-\rho)^2\hat{\sigma}_{ipt}^2 + (1-\rho^2)\sigma_\varepsilon^2$, which proves (4.43). The pair $(\hat{\theta}_{ip,t+1}, \log \theta_{ip,t+1})$ is jointly normal, with posterior-mean variance $\hat{\sigma}_{ipt}^2 - \hat{\sigma}_{ip,t+1}^2$ and covariance $(1-\rho)\hat{\sigma}_{ipt}^2$. At tenure zero, both means equal θ_0 , and the posterior-mean variance, covariance, and demand variance are $\sigma_0^2 - \hat{\sigma}_{ip,t+1}^2$, σ_0^2 , and $\sigma_0^2 + \sigma_\varepsilon^2$. These moments define the transition for the Bellman equations on the state $(\hat{\theta}_{ipt}, \log \theta_{ipt}, \tau_{ipt}; a_{it}, \psi_{it}, \varphi_i, m_i)$. Both continuous coordinates are needed, because equal demand terms need not imply equal continuation values. Online Appendix OA.B.1 gives the complete system.

B.2 Mapping Shipment Fixed Effects to Expertise

Under full information about permanent expertise, its loading in $\log \mathcal{B}$ is $(\eta-1)/\eta$. In (2.11), its total loading in $\log$ real shipments is therefore $(\eta-1)^2/\eta + (\eta-1)/\eta = \eta-1$. After productivity and aggregate components are removed, this gives the benchmark mapping $\alpha_{ip}^{PV} \simeq (\eta-1) \log \theta_{ip}$. Demand persistence changes shipment dynamics but not this loading. We use this mapping as a benchmark for interpreting the estimated fixed effects; see Online Appendix OA.B.2.

C Restricted Planner: Formulation and Proof Sketches

Given the initial firm distribution, inherited capabilities, and predetermined quantities, the restricted planner maximizes $W^{SP} \equiv \sum_{t \geq 0} \beta^t \log \bar{C}_t$ over physically feasible production, product switching, startup operation, entry, and pre-entry investments. It uses firms' infor-

mation, allocates production efficiently given that information, and keeps the CE valuation rules. From age two, each line is valued independently of its effects on firm exit and other lines and receives its vacant-line value at endogenous exit. At age one, only the initial line receives this value; initially vacant lines receive zero. Startup operation and entry use the same values, which create no resources. Capability technologies and costs are those in Section 2.4; private optimality of capability investment is not imposed. The CE information structure, timing, and physical laws of motion apply, with symmetric markets, positive vacant-line values, and a common rule for breaking ties. Online Appendix OA.C.5.1 gives the formal definition.

C.1 Conditional Production and Verification

Let $R_t = \bar{L} - \mathcal{F}_t^o - \mathcal{F}_t^s - \mathcal{F}_t^e > 0$ be labor available for variable production and $S_t = \int \mathcal{B}(\xi)^\eta z_i^{\eta-1} \mu_{p,t-1}(d\xi) \in (0, \infty)$ the productive index of a symmetric market. Conditional on planning information, the law of large numbers and Hölder's inequality give the optimal allocation whenever quantities can be chosen:

$$\ell_{pt}(\xi) = \frac{R_t}{N} \frac{\mathcal{B}(\xi)^\eta z_i^{\eta-1}}{S_t}, \quad C_{pt} = \frac{R_t}{N} S_t^{1/(\eta-1)}. \quad (\text{C.1})$$

Holding search and entry resources fixed, an active line's direct marginal contribution in labor units is $\ell_{pt}(\xi)/(\eta - 1) - f^o$, equal to CE expected profit divided by the wage. Online Appendix OA.C.5.2 gives the derivation.

The planner's decision rules discount labor returns at $\beta_t^{SP} = \beta R_t/R_{t+1}$ and use the CE line and startup Bellman equations with these operating returns and assigned exit values. Attempted entry $\mathcal{E}_t \geq 0$ satisfies $\mathbb{E}_\varphi V_t^e(\varphi) \leq f_t^e$, with equality if $\mathcal{E}_t > 0$, and uses $f_t^e \mathcal{E}_t$ units of labor. Search and production follow actual firm continuation.

Lemma C.1 (Verification). *Fix an aggregate path satisfying the measurability, integrability, and uniform payoff-tail conditions in Online Appendix OA.C.5.4. The product-line and startup-operation Bellman equations have a finite measurable value function attained by a measurable policy, unique among solutions with vanishing discounted expected terminal values under every admissible policy.*

Proof sketch. The uniform tail condition gives convergence of finite-horizon values. Finite action sets admit measurable maximizers; iterating the Bellman inequality and applying the terminal condition gives attainment and uniqueness. $\square$

For welfare, we also require $\sum_{t \geq 0} \beta^t |\log \bar{C}_t| < \infty$ and assume that the planner's maximum is attained.

C.2 Replication and the Value of Capability Choice

Proposition C.1 (Replication). *Fix (m_0, z_0) and their acquisition costs. Under the maintained restrictions and regularity conditions, the restricted planner's admissible stationary allocations coincide with stationary CE allocations, retaining startup operation and free entry but not private capability-investment optimality. Under a common equilibrium selection, consumption and welfare coincide.*

Proof sketch. Conditional production gives the CE allocation and operating profit divided by the wage. Dividing the CE decision and entry conditions by the constant wage yields the planner conditions. Lemma C.1 and common tie breaking give identical values and policies. Common physical transitions and resource constraints support the same stationary distribution. Conversely, CES prices in (2.9) and (2.36) support any admissible stationary planner allocation; multiplying planner values by the wage recovers the CE decision and entry conditions. Online Appendix OA.C.5.5 gives the full proof and transition extension. $\square$

Corollary C.1 (Value of Capability Choice). *Starting from a stationary CE's firm distribution, inherited capabilities, and predetermined quantities, allowing pre-entry capability choice weakly raises welfare: $W^{SP} \geq W^{CE}$.*

Proof sketch. Setting $a_t^m = m_0^{CE}$ and $a_t^z = z_0^{CE}$ at every date reproduces the CE capabilities and costs. The stationary CE remains feasible for the planner, implying $W^{SP} \geq W^{CE}$. $\square$