Winners and Losers: Competition, Creative Destruction, and Labor Income Risk*

Brice Green[1], Leonid Kogan[1,3], Dimitris Papanikolaou[2,3], and Lawrence D. W. Schmidt[1]

[1] *MIT Sloan School of Management*
[2] *Kellogg School of Management*
[3] *National Bureau of Economic Research*

March 2026

Abstract

Using U.S. administrative data, we find that technology-driven creative destruction in the product market passes through to worker earnings. The passthrough to incumbent worker earnings is both asymmetric and concentrated: profit drops from rival innovations lead to proportionally greater earning declines and changes in the likelihood of job destruction than profit gains from their own firm's innovations, while top workers are significantly more exposed than the average worker. We develop an endogenous-growth model with monopsonistic labor markets and worker heterogeneity that replicates this asymmetry and the distribution of earnings risk. Creative destruction exposes high-income workers to concentrated downside risk while increasing upward mobility for lower-income workers, shaping the welfare consequences of innovation policy.

*This paper subsumes the earlier paper titled "Technological Innovation and Labor Income Risk", by Kogan, Papanikolaou, Schmidt, and Song (2021), which was performed pursuant to a grant from the U.S. Social Security Administration (SSA) funded as part of the Retirement and Disability Research Consortium. The authors also gratefully acknowledge financial support from the Becker-Friedman Institute for Research in Economics and the MIT Sloan School of Management. We thank Nicholas Bloom, Carter Braxton, Stephane Bonhomme, Gregor Jarosch, Chad Jones, Pete Klenow, Thibaut Lamadon, Danial Lashkari, Huiyu Li, Ilse Lindenlaub, Elena Manresa, Derek Neal, Serdar Ozkan, Luigi Pistaferri, Jae Song, Christopher Tonetti, Toni Whited, and participants in various seminars and conferences helpful comments and discussions. Further, we thank Fatih Guvenen, Serdar Ozkan, and Jae Song for sharing their replication code. Brandon Dixon, Luxi Han, Alejandro Hoyos Suarez, Kyle Kost, Andrei Nagornyi, Bryan Seegmiller, and Jiaheng Yu provided outstanding research assistance. 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 (Data Management System (DMS) number: P-7503840, Disclosure Review Board (DRB) approval numbers: CBDRB-FY23-SEHSD003-046, CBDRB-FY23-SEHSD003-073, CBDRB-FY25-0100.

The process of economic growth is inseparable from creative destruction, in which new entrants displace incumbent firms (Aghion and Howitt, 1992). In canonical macroeconomic models, creative destruction is largely benign: investors diversify firm-specific risks, and workers move costlessly from declining firms to expanding ones. In practice, a significant fraction of a worker's human capital is specific to a particular job, hence workers view these jobs as imperfect substitutes. As a result of firms internalizing their monopsony power, worker earnings often move with the fortunes of the employer. In this paper, we document the extent to which workers share in the gains and losses from creative destruction and interpret these facts through the lens of a structural model. A central insight is that creative destruction is a powerful force for earnings redistribution: it increases upward mobility for lower-paid workers while exposing higher-paid workers to substantial downside risk.

We begin by showing that increases in a firm's own innovation raise future profits and the total pay of its workers, while innovation by competing firms lowers both. At the firm level, these responses appear largely symmetric: firms that successfully innovate expand employment and their wage bill, while firms that lose product lines contract. However, this symmetry breaks down when we shift the focus from firm aggregates to individual workers. Using the panel structure of administrative employer–employee matched data from the United States, we estimate how individual earnings respond to innovation by the worker's employer and its competitors. Workers in innovating firms experience earnings gains, while workers in firms whose competitors innovate suffer losses—but these responses are strikingly asymmetric. The passthrough from profits to earnings is about 0.14 for own-firm innovations, consistent with existing estimates of profit sharing. For competitor innovations, the implied passthrough is 0.29—more than twice as large—even though the corresponding decline in firm profits is similar in magnitude.

This asymmetry in passthrough highlights the unique role of innovation in driving firm profits and worker earnings. When a firm innovates and expands, it raises wages and hires additional workers; part of the increase in the wage bill therefore accrues to new hires, not incumbents. Incumbent workers are already employed and do not share in the gains from employment growth. By contrast, when a competitor innovates and the firm loses market share, both the wage decline and the associated job loss fall squarely on incumbents. Incumbent workers thus bear the full cost of creative destruction but capture only a fraction of its benefits. We provide three additional pieces of evidence consistent with this interpretation. First, we find that increases in competitor innovation significantly raise the probability of job destruction for incumbent workers—measured using a proxy for involuntary job loss based on whether a worker enters nonemployment or switches employers while simultaneously experiencing a sharp decline in earnings. Second, estimates from quantile regressions reveal that, own-firm innovation produces positively skewed earnings growth—concentrating gains among a small subset of workers—whereas competitor innovation yields negatively skewed earnings changes, with most workers experiencing modest declines and a minority suffering large losses. Last,

we show that innovations that are more distinct relative to the firm's existing technological base generate large increases in profits, payroll, and employment but only modest earnings gains for incumbent workers relative to equally valuable, less novel innovations.

Both the gains and the losses from innovation are disproportionately concentrated at the top of the earnings distribution. Workers in the top 5% of their industry's earnings distribution see earnings gains of 2.6 percentage points following a one-standard-deviation increase in own-firm innovation, compared to a 4 percentage point earnings decline from an equivalent increase in competitor innovation. The implied passthrough coefficients for these top earners are two to three times larger than for the average worker, consistent with higher-skilled workers facing less elastic labor supply. This pattern also varies systematically across industries: the passthrough of innovation to worker earnings is significantly larger in industries with high rates of innovative activity (and potentially more scope for creative destruction) and in industries in which firms likely face less elastic labor supply curves: those with greater within-industry income dispersion as well as those with higher average wages. Once again, these exposures are highest among the highest-paid workers in these industries.

We next interpret these estimates through the lens of a model of creative destruction and labor market frictions. The model combines several key elements from the literature. Firms compete in producing different varieties, as in Aghion and Howitt (1992). As in Peters (2020), firms can engage either in internal or external innovation; internal innovation increases the productivity of existing product lines while external innovation allows the firm to steal business from a competitor; a competitive fringe of potential entrants provides an additional source of creative destruction. Unlike Peters (2020), the rate of industry innovation is stochastic: the model features innovation booms in the spirit of Berlingieri, Ridder, Lashkari, and Rigo (2025). The leading firm produces output using two complementary labor inputs: laborers and managers. The key distinction between laborers and managers in production is that the optimal ratio of managers to laborers is increasing in firm size.

Workers, either laborers or managers, optimally choose their preferred job out of a finite set; jobs differ in their match-specific productivity as in Roy (1951). Firms internalize their labor market power, since jobs are imperfect substitutes; hence they pay workers below their marginal product of labor as in Gouin-Bonenfant (2022); Seegmiller (2024). Laborers have a more elastic labor supply than managers; since the more productive firms need to hire more workers but also a higher ratio of managers to laborers, they face more inelastic labor supply curves. Consequently, the more productive firms have higher profit margins as they optimally mark down wages further below the marginal product of labor.

In sum, incumbent workers in the model face two forms of income risk linked to innovation. The first operates through wages in continuing jobs: internal innovation raises a firm's relative productivity and increases labor demand and wages, while competitor innovation has the opposite

effect. This channel generates roughly symmetric responses to own-firm and competitor shocks. The second source of risk—creative destruction—drives the asymmetry we document in the data. A worker's human capital is tied to the product line on which she is employed. When a competitor or entrant displaces the incumbent producer, the job disappears and the worker must draw a new match from her remaining options, all of which pay less. Because high-income workers sort into matches where their human capital is most valuable, displacement imposes disproportionately large losses on them. This mechanism also limits the degree to which incumbent workers benefit from their own firm's innovation. Workers gain when the product line they work on improves, but they do not share in gains that arise from the firm expanding into other lines, such as when it steals a product from a competitor. As a result, incumbent workers capture only a fraction of the upside from innovation but bear most of the downside.

The calibrated model quantitatively replicates the key relation between technological progress and earnings risk that we uncover. In addition, the model reproduces key empirical features of worker-level earnings risk that are not explicit calibration targets. As in Guvenen, Karahan, Ozkan, and Song (2021), high-income workers face disproportionately negative skewness, reflecting their exposure to creative destruction when high-match-quality jobs disappear. The model reproduces these non-Gaussian features of the earnings distribution without targeting them in our calibration.

We conclude by performing two quantitative experiments. First, we explore implications of the well-documented secular decline in business dynamism (Haltiwanger, 2012; Akcigit and Ates, 2019) through the lens of the model. We show a related fact: the value share of radical innovation—patents that are technologically distinct from a firm's existing base—fell from roughly 40 percent in the 1990s to 20 percent in the 2010s. Recalibrating the model to match this shift, we find that the decline in creative destruction can simultaneously account for several prominent macroeconomic trends: a falling aggregate labor share, rising firm concentration, declining dispersion in firm growth rates, and—consistent with the "Great Micro Moderation" (Sabelhaus and Song, 2010; Salgado, Bloom, Guvenen, Pistaferri, and Sabelhaus, 2023)—a reduction in the cross-sectional dispersion of firm-level employment growth and individual earnings growth. A key implication is that declining creative destruction lowers the earnings risk that workers face but at the same time increases inequality, as reduced displacement allows productivity differences across firms and workers to persist and widen.

Last, we use the calibrated model to study the welfare consequences of subsidizing innovation, or, equivalently, to exogenous shifts in the cost of innovation. In contrast to canonical growth models in which innovation unambiguously raises welfare, our environment features incomplete markets and uninsurable labor-income risk. As a result, incumbent workers value innovation very differently depending on how it affects creative destruction. A temporary subsidy that raises internal innovation generates welfare gains close to the complete-markets benchmark. By contrast, a subsidy that increases external innovation and entry—thereby raising the rate of creative destruction—reduces

welfare for most workers, especially those at the top of the earnings distribution most exposed to displacement risk. High-income workers would pay several percent of annual consumption to avoid such a policy, while lower-paid workers benefit from the increased chance of drawing a better match.

These results highlight how increased business dynamism can redistribute earnings even as it reduces welfare for some workers, echoing the intuition of Krebs (2007) that the welfare costs of business cycles are higher when incorporating job displacement risk. That said, these calculations adopt the ex-post perspective of current incumbent workers who already know current labor earnings, which is the relevant one when considering changes in policy along a transition path. Intriguingly, behind the veil of ignorance (i.e., before learning one's productivity), these perceived costs vanish and each agent's willingness to pay coincides with the complete markets benchmark. Hence, our results also point to an additional channel—income risk—through which resistance to technological change can arise, especially among higher-paid workers.

In sum, our work emphasizes the role of creative destruction as a driver of worker earnings risk. Our paper combines insights from Schumpeterian endogenous growth models (Aghion and Howitt, 1992; Klette and Kortum, 2004; Lentz and Mortensen, 2008; Peters, 2020; Acemoglu, Akcigit, Alp, Bloom, and Kerr, 2018; Jones and Kim, 2018) with models of worker comparative advantage (Roy, 1951; Lagakos and Waugh, 2013; Hsieh, Hurst, Jones, and Klenow, 2019; Burstein, Morales, and Vogel, 2019; Caunedo, Jaume, and Keller, 2023; Galle, Rodríguez-Clare, and Yi, 2023). Further, our work contributes to a nascent literature that emphasizes the importance of creative destruction for firm growth in the data (Akcigit and Kerr, 2018; Argente, Lee, and Moreira, 2024; Berlingieri et al., 2025), but differs in placing individual workers—rather than firms—at the center of the analysis. Specifically, models with creative destruction typically assume frictionless labor markets, which imply that creative destruction does not affect worker earnings. Two notable exceptions are Aghion, Akcigit, Bergeaud, Blundell, and Hemous (2018), who study the impact of innovation on income inequality, and Aghion, Akcigit, Deaton, and Roulet (2016), who show that the degree of creative destruction is negatively related to measures of subjective well-being.

Our work is also related to a sizeable literature in labor economics that estimates the passthrough of firm shocks to workers (Guiso, Pistaferri, and Schivardi, 2005; Lagakos and Ordoñez, 2011; Fagereng, Guiso, and Pistaferri, 2018; Friedrich, Laun, Meghir, and Pistaferri, 2019; Lamadon, Mogstad, and Setzler, 2022; Chan, Xu, and Salgado, 2023). That literature interprets passthrough estimates as informative about the degree of insurance provided by firms to workers. By contrast, in our model, the passthrough arises because workers differ in their match-specific productivity, which provides firms with labor market power. That said, our empirical estimates imply average passthrough of firm profits to worker earnings that is in line with the literature (Card, Cardoso, Heining, and Kline, 2018; Lamadon et al., 2022). In addition, the fact that higher-paid workers have higher exposure to innovation by their employer echoes the results in Friedrich et al. (2019), who

find that the passthrough of firm-specific permanent productivity shocks is mostly concentrated among high-skilled workers. Some of these papers focus on the impact of firm innovation on the earnings of individual workers (van Reenen, 1996; Aghion, Bergeaud, Blundell, and Griffith, 2017; Kline, Petkova, Williams, and Zidar, 2019; Toivanen and Väänänen, 2012; Bell, Chetty, Jaravel, Petkova, and Van Reenen, 2019; Akcigit, Grigsby, and Nicholas, 2017; Aghion, Akcigit, Hyytinen, and Toivanen, 2018). However, none of these papers focus on how worker earnings respond to innovation by their firm's competitors and are therefore silent on the link between the overall level of creative destruction and worker earnings risk.

Finally, our evidence has implications for asset prices. Because creative destruction concentrates earnings risk among high earners, who hold a disproportionate share of aggregate wealth, firms more exposed to creative destruction are likely to lose value in high marginal utility states, providing an additional microfoundation for the return spread between value and growth stocks (Kogan, Papanikolaou, and Stoffman, 2020). The presence of uninsurable displacement risk also raises the precautionary savings motive for high earners, connecting to recent work on the global savings glut of the rich (Mian, Straub, and Sufi, 2021; Braxton, Herkenhoff, Huang, Nattinger, Rothbaum, and Schmidt, 2025).

1 A Motivating Framework

To motivate the empirical analysis, we begin with a stripped-down version of the model in Section 3. The goal is to isolate the core mechanisms without the algebraic detail. This simplified environment does two things. First, it illustrates how labor-market frictions transmit creative destruction in product markets into workers' earnings. Second, it delivers a transparent estimating equation that guides our empirical specification. The full derivation and additional structure are provided in Appendix C.1.

Aggregate output $Y(t)$ is produced by a perfectly competitive final-goods sector that aggregates a continuum of industry outputs $Y(j,t)$ via a Cobb-Douglas aggregator across industries, so the elasticity of substitution across industries is one and spending is equalized across industries.

Within each industry j, output is a CES composite of a continuum of intermediate varieties indexed by v:

$$Y(j,t) = \left(\int_0^1 Y_j(v,t)^{\frac{\sigma-1}{\sigma}} dv \right)^{\frac{\sigma}{\sigma-1}}. \tag{1}$$

where $\sigma > 1$ is the elasticity of substitution across varieties within an industry.

For simplicity, we assume for now that each firm produces a single intermediate variety, so the index v labels both firms and varieties. In the full model in Section 3, we relax this assumption and allow firms to supply multiple varieties within an industry. An intermediate-good producer

combines its current technology level $Q_j(v,t)$ with labor to produce output:

$$Y_j(v,t) = Q_j(v,t)\, l(w_j(v,t)). \tag{2}$$

Because firms compete monopolistically within each industry, each variety producer sets price as a constant markup over marginal cost. This markup reflects the elasticity of substitution σ that governs within-industry competition.

In addition to their market power on the product side, firms also possess monopsony power in the labor market. Following Lamadon et al. (2022), each intermediate-good producer faces an upward-sloping labor supply curve for worker services,

$$l^s(w) = w^{\frac{1}{\varepsilon}} \tag{3}$$

where $1/\varepsilon$ is the elasticity of labor supply and w is the posted wage of each firm—relative to total output $Y(t)$. Hence, a worker's per-period earnings are equal to $Y(t)\, w_j(v,t)$.

The marginal cost of labor is $w\,(1+\varepsilon)$, so the ratio of marginal to average labor cost is simply $1+\varepsilon$. For simplicity, we interpret the labor supply curve (3) as applying to the extensive margin: each worker supplies one unit of efficiency labor. The full model in Section 3 generalizes this environment to allow for both intensive and extensive margins of labor supply.

Firms choose their labor input to maximize profits, taking into account their market power on both the product and labor sides. Because each firm is a monopolistic competitor in its variety and a monopsonist in its local labor market, it internalizes how its production decision affects both the price it charges and the wage it must pay. Combining these two margins, the firm's total revenue over average cost satisfies

$$\frac{P_j(v,t)\, Y_j(v,t)}{Y(t)\, w_j(v,t)\, l_j(v,t)} = \frac{\sigma}{\sigma-1} \times [1+\varepsilon]. \tag{4}$$

That is, profit margins depend on the standard CES markup term (which decreases with the elasticity of substitutions across varieties σ) and a monopsony adjustment term (which increases with the inverse elasticity of labor supply ε). Given this wedge, the firm's labor demand satisfies

$$l_j(v,t) = \left(\frac{\bar{A}}{\hat{Q}_j(v,t)}\right)^{1-\sigma} \left[\frac{\sigma}{\sigma-1} w_j(v,t)(1+\varepsilon)\right]^{-\sigma}. \tag{5}$$

In equilibrium, wages adjust so that firm-level labor demand is equal to

$$l_j(v,t) = \left(\frac{\bar{A}}{\hat{Q}_j(v,t)}\right)^{\frac{1-\sigma}{1+\varepsilon\sigma}} \left[\frac{\sigma}{\sigma-1}(1+\varepsilon)\right]^{-\frac{\sigma}{1+\varepsilon\sigma}} \tag{6}$$

which depends only on the firm's productivity relative to its industry average,

$$\hat{Q}_j(v,t) \equiv \frac{Q_j(v,t)}{\bar{Q}_j(t)}, \qquad \text{where} \quad \bar{Q}_j(t) \equiv \int_0^1 Q_j(v,t)dv, \tag{7}$$

and where $\bar{A}$ is a constant derived from the CES demand system that captures the level of aggregate productivity.

The key implication of monopsony power is that changes in a firm's relative productivity $\hat{Q}_j(v,t)$ now affect both wages and employment. Under perfect competition ($\varepsilon \downarrow 0$), a productivity shift alters only the quantity of labor demanded; the wage is pinned down by the common market wage. With monopsony power ($\varepsilon > 0$), a more productive firm raises both its wage and employment. As labor supply becomes more inelastic (larger ε), the adjustment occurs increasingly through wages rather than quantities.

A central focus of our analysis is how changes in firm productivity shape the wage dynamics of incumbent workers. Following an innovation shock that raises productivity for a subset of varieties in industry j, firm profits evolve according to

$$\mathrm{d}\log \pi_j(v,t) = \left(1+\frac{1}{\varepsilon}\right)\frac{\sigma-1}{\sigma+\frac{1}{\varepsilon}}\left(\mathrm{d}\log Q_j(v,t) - \mathrm{d}\log \bar{Q}_j(t)\right). \tag{8}$$

This expression highlights that future profits depend on the change in the firm's technology *relative* to its competitors, $\hat{Q}_j(v,t)$. Innovations that disproportionately benefit a subset of firms raise their relative productivity, while simultaneously lowering the relative position of competitors whose productivities do not change. In this sense, when the gains from innovation are concentrated, industry-level technological progress produces both winners and losers.

Under monopsonistic labor markets, these relative-productivity shocks affect not only profits but also the earnings of incumbent workers who remain at the firm. The equilibrium evolution of their wages is

$$\mathrm{d}\log w_j(v,t) = \frac{\sigma-1}{\sigma+\frac{1}{\varepsilon}}\left(\mathrm{d}\log Q_j(v,t) - \mathrm{d}\log \bar{Q}_j(t)\right) \quad = \left(1+\frac{1}{\varepsilon}\right)^{-1}\mathrm{d}\log \pi_j(v,t). \tag{9}$$

The elasticity of worker earnings with respect to innovation-driven changes in profits is decreasing in ε: wages respond more strongly when labor supply is less elastic.

Importantly, worker earnings are sensitive not only to improvements in their own firm's technology (consistent with the evidence in van Reenen, 1996; Kline et al., 2019), but also to innovations by competitors. This simple environment captures a diffuse form of creative destruction, a hallmark of Schumpeterian models: holding own-firm fundamentals constant, increases in the average productivity of rival firms reduce market share and depress labor demand. In what follows, we construct empirical measures of competitor innovation and provide new evidence linking these shocks to the earnings of incumbent workers.

In sum, the simple model delivers two testable predictions. First, worker earnings respond to changes in their employer's productivity *relative* to competitors: own-firm innovation raises earnings, while competitor innovation lowers them. Second, the passthrough coefficient—$(1+1/\varepsilon)^{-1}$—increases when labor supply is more inelastic, so workers with less elastic labor supply should exhibit greater

exposure to innovation shocks.

The full model in Section 3 extends this framework in three directions. First, it endogenizes innovation in a dynamic environment. Second, it microfounds labor supply using a tractable Roy (1951) model in which workers reallocate across firms. Third—and most importantly for generating the asymmetric passthrough we find in the data—it incorporates *business stealing*: innovations by competitors can cause a firm to lose its product line entirely. Because a worker's human capital is tied to a specific product line, business stealing destroys valuable matches and imposes large, concentrated losses on displaced workers.

2 Creative Destruction, Firm Profits, and Worker Wages

We now document that innovation-driven creative destruction in the product market passes through to incumbent workers' earnings. The passthrough is both asymmetric and concentrated: profit losses from rival innovations translate into proportionally larger earnings declines than the gains from own-firm innovation, and these effects fall disproportionately on high-income workers. Our empirical strategy proceeds in two steps. First, we construct direct proxies for shifts in a firm's own technology and in the technology of its competitors, guided by equations (8) and (9). Second, we estimate how these innovation shocks affect firm profits and the extent to which these effects pass through to the wage earnings of incumbent workers.

2.1 Data and Measurement

We begin by briefly summarizing the data on labor income and firm innovation outcomes used in our analysis. All details are relegated to Appendix B.1.

Firm Innovation

Our first step is to build empirical measures of innovation at the firm and competitor level. We follow Kogan, Papanikolaou, Seru, and Stoffman (2017) (henceforth KPSS), who estimate the market value of a patent from the stock price reaction of the innovating firm at the time of grant. Their approach delivers a direct, forward-looking measure of patent value, but it is available only for public firms because it relies on observed stock returns. We refer to this as the market value of a patent.

For each firm f in year t, we define:

$$A_{f,t} = \frac{\sum_{j \in P_{f,t}} KPSS_j}{K_{ft}}. \tag{10}$$

Here, $KPSS_j$ is the KPSS market value of patent j, and $P_{f,t}$ is the set of patents filed by firm f in year t. The numerator is the total market value of a firm's patents filed in year t. The denominator scales this by firm size, using the firm's book assets K_{ft} as our baseline. This adjustment accounts for the fact that larger firms file more patents. Results are robust if we instead scale by enterprise

value. In the context of our model, $A_{f,t}$ proxies for changes in the quality of firm f's production technology, $Q_j(v,t)$.

We also construct an industry-level measure of competitor innovation:

$$A_{I\setminus f,t} = \frac{\sum_{f' \in I\setminus f} \sum_{j \in P_{f',t}} KPSS_j}{\sum_{f' \in I\setminus f} K_{f't}}. \tag{11}$$

This is the ratio of the total patent value of other firms in the same 3-digit SIC industry (excluding firm f) to the sum of their book assets.

The two measures differ in timing: for own-firm innovation ($A_{f,t}$), we date patents by their application year to avoid capturing anticipatory wage adjustments before the grant date; for competitors' innovation ($A_{I\setminus f,t}$), we use the grant year. This choice is not the main driver of our findings, as results are qualitatively similar if we also date the firm's patents as of their grant date. Further details on the timing convention are provided in Appendix B.1.

Firm Wage Bill and Employment

We measure firm wages and employment from the Longitudinal Business Database (LBD), linked to Compustat through Employer Identification Numbers (EINs). The sample runs from 1976 to 2014 and covers 12,500 firms with 132,000 firm-year observations. These are large, publicly traded firms. The average wage bill is $282 million, compared to a median of $27 million. The average firm-year employs 8,200 workers, while the median employs 930. Both pay and employment are highly skewed.

Worker earnings

We measure worker earnings using the CPS merged with Social Security's Detailed Earnings Record (DER). The CPS/DER panel tracks individuals across time and across firms from 1980 to 2020. These data report W-2 earnings and employer identifiers, which we use to link workers to Compustat firms by Employer Identification Numbers (EINs). We focus on prime-age workers (25–55), exclude the self-employed, and drop individuals earning less than a threshold equal to 20 hours per week at the federal minimum wage for 13 weeks (Guvenen, Ozkan, and Song, 2014). After requiring sufficient income history and firm controls (e.g. lagged R&D), we obtain a worker panel with 9,700 firms, 780,000 people, and 5.6 million person-years. The corresponding firm panel has 12,500 firms and 132,000 firm-years. These worker records give us the micro-level earnings outcomes we need to evaluate the effects of firm and competitor innovation defined above.

Our main outcome is the cumulative growth rate of worker earnings over h years:

$$g_{i,t:t+h} \equiv w^i_{t+1,t+h} - w^i_{t-2,t}. \tag{12}$$

For the bulk of our analysis we will focus on 5 year horizons, $h = 5$. To remove life-cycle effects, we

follow Autor, Dorn, Hanson, and Song (2014) and compute adjusted average earnings:

$$w^i_{t,t+h} \equiv \log\left(\frac{\sum_{j=0}^{h} \text{W-2 earnings}_{i,t+j}}{\sum_{j=0}^{k} D(\text{age}_{i,t+j})}\right). \tag{13}$$

Here, W-2 earnings are summed across all jobs for person i in year t, and $D(\text{age})$ adjusts for the mean life-cycle profile in earnings (Guvenen et al., 2014).

Two points are worth noting. First, averaging income growth over multiple years, as in (12)–(13), emphasizes persistent changes and smooths over transitory shocks or unemployment spells. Our baseline compares forward five-year earnings with the prior three years; other horizons give similar patterns. Second, the panel allows us to follow workers even if they leave a Compustat firm. A worker enters the sample in year t if employed by a matched firm, but subsequent earnings include income from any new employer. This ensures that $g_{i,t:t+h}$ captures the worker's true realized earnings growth, not just the growth in earnings during her tenure with a given firm.

2.2 Innovation, Profits, and Wage Earnings

We now estimate the empirical counterparts of equations (8) and (9), which relate profit and wage growth to firm and competitor innovation.

Firms

We begin by estimating the link between innovation, profits and wage earnings at the firm level. To match the worker-level analysis below, we measure cumulative growth in profits and average wages over horizon h.

$$g_{f,t:t+h} \equiv \log\left[\frac{1}{|h|}\sum_{\tau=1}^{h} X_{f,t+\tau}\right] - \log X_{f,t}. \tag{14}$$

We then estimate the following specification,

$$g_{f,t:t+h} = a\, A_{f,t} + b\, A_{I\setminus f,t} + c\, \mathbf{Z}_{ft} + u_{ft}. \tag{15}$$

The coefficients a and b capture how firm and competitor innovation affect profitability and average wages at the firm level. The controls $\mathbf{Z}_{ft}$ include firm size (log market capitalization and log book assets), R&D intensity, the KPSS value of the firm's patent portfolio (assuming 20% depreciation), the firm's idiosyncratic volatility, lagged gross profitability, and SIC 3-digit industry and year fixed effects. Standard errors are clustered by industry and year. To compare magnitudes, we normalize both $A_{f,t}$ and $A_{I\setminus f,t}$ to unit standard deviation. Appendix B.1 contains further details.

The top two rows of Figure 1 plot the responses of profits and average wages to firm and competitor innovation. The pattern is clear. When firms innovate, both profits and wages rise. A one-standard-deviation increase in firm innovation raises profits by 5.6 percent and average wages by 2.0 percent over the next five years. Competitor innovation works in the opposite direction. A one-standard-deviation increase in competitor innovation reduces focal firm profits by 4.1 percent,

with only a small and statistically insignificant effect on worker pay.

At first glance, the flat response of average wages to competitor innovation might suggest that competitor innovation does not pass through to workers. That conclusion is misleading. It ignores the extensive margin: firms adjust employment as well as wages. The bottom two rows of Figure 1 plot the responses of employment and the total wage bill. We see that firm employment strongly responds to innovation: innovation by the firm leads to a 2.9 percent increase in employment and a 4.9 percent increase in the wage bill over the same period. Similarly, innovation by the firm's competitors is followed by a 2.7 percent decline in employment and a 3.4 percent decline in the wage bill over the next five years.

A final point is the absence of pre-trends in these firm-level responses. Figure 1 plots the estimated coefficients a and b across horizons $h = -5 \dots 10$. For $h = -1$ to $h = -5$, both coefficients are small and statistically insignificant. In other words, past changes in profits, wages, or employment show no correlation with either A_f or $A_{I \setminus f}$. This lack of pre-existing trends reduces concerns that our innovation measures simply capture unobserved heterogeneity driving firm outcomes.

Individual Workers

A limitation of the firm-level analysis for studying worker outcomes is that it conditions on continuing employment. Individual workers can still experience earnings losses if they lose their jobs—even as average wages rise at the firm level, for example, if firms selectively lay off lower-paid employees. We therefore move from firm aggregates to individual earnings to see how innovation shocks transmit directly to workers.

We estimate:

$$g_{i,t:t+h} = a\, A_{f,t} + b\, A_{I \setminus f,t} + c\, \mathbf{Z}_{ft} + d\, \mathbf{X}_{i,t} + \varepsilon_{i,t}, \tag{16}$$

where $g_{i,t:t+h}$ is the cumulative growth of worker i's earnings. A positive a means workers share in their firm's innovation rents, while a negative b means competitor innovation erodes their earnings. These coefficients are identified by comparing the future earnings path for workers employed in firms that differ in their degree of innovation exposure but are otherwise similar in terms of age and current and past level of wage earnings. We control for the same set of firm-level variables $\mathbf{Z}_{ft}$ as in equation (15), plus a vector of worker controls $\mathbf{X}_{i,t}$ that includes: non-parametric controls for worker age, past earnings level, and recent earnings growth rates. We also include worker income $\times$ industry and worker income $\times$ year fixed effects, with workers partitioned into five bins of the within-industry, within-year earnings distribution. We weight observations by the inverse of the number of workers in each firm-year and cluster standard errors at the industry-year level. Appendix B.1 contains further details.

Figure 2 reports the estimated a and b coefficients for incumbent workers over horizons of three to ten years. The results confirm that innovation shocks pass through to individual workers.

Firm innovation raises worker earnings, while competitor innovation lowers them—and the latter effect is larger and more persistent. At the five-year horizon, a one-standard-deviation increase in firm innovation (A_f) raises cumulative earnings growth by 0.8 percentage points. The same increase in competitor innovation ($A_{I\setminus f}$) reduces cumulative earnings growth by 1.2 percentage points. Looking across horizons, competitor innovation continues to depress earnings long after the shock. A one-standard-deviation increase cuts earnings by 1.1 percent at three years and 1.2 percent at ten years. By contrast, the effect of own-firm innovation fades: earnings gains are 0.9 percent at three years and only 0.6 percent at ten years. Thus, workers benefit modestly from their own firm's innovations, but competitor innovations impose larger and more durable earnings losses.

Contrasting the estimates from the worker-level to the firm-level regressions yields two insights. First, the ratio of changes in worker earnings to changes in firm profits gives the passthrough coefficient—equation (9) in the simple model. Our estimates of a from equations (16) and (15) imply a passthrough of about $0.8/5.7 \approx 0.14$, which is close to values in the literature (e.g., Card et al., 2018; Lamadon et al., 2022). By contrast, the estimates of b imply a much larger worker earnings decline relative to firm profits: $1.2/4.1 \approx 0.3$. Put differently, for every 1 percent fall in firm profits from competitor innovation, worker earnings drop by about 0.3 percent. The passthrough of innovation shocks is thus asymmetric: smaller for profit gains from own innovation, larger for losses from competitor innovation.

Second, incumbent workers bear a disproportionate share of the costs of creative destruction. Comparing Figure 2 (earnings of incumbent workers) with the second row of Figure 1 (average pay for all continuing employees) reveals that incumbents' earnings fall significantly more than the firm-wide average wage following competitor innovation. The gap reflects worker turnover: when a firm expands after innovating, part of the surplus accrues to new hires; when a firm contracts after competitors innovate, incumbents absorb the full burden. Incumbents thus capture fewer of the gains from own-firm innovation and bear more of the losses from competitor innovation.

Worker Heterogeneity

We next ask which workers gain most from their own firm's innovation and which are most exposed to competitor-driven losses. The simple model in Section 1 predicts that the passthrough of profit shocks depends on the elasticity of labor supply, ε. A natural starting point is that higher-skilled workers face lower labor supply elasticities than the average worker. To test this, we re-estimate equation (16), allowing the coefficients a and b to vary with a worker's prior income relative to others in the same industry and year. Workers are grouped into five bins of the within-industry, within-year earnings distribution: bottom quartile (0–25th percentile), second quartile (25–50th), third quartile (50–75th), upper decile excluding the top (75–95th), and top ventile (95–100th). The coefficients on the control variables are also allowed to vary by income group.

Figure 4 reports the resulting a and b estimates for each group across horizons of three to ten

years, together with the implied profit-sharing elasticities. The gradient is steep: higher-paid workers are much more exposed to innovation shocks than lower-paid workers. At the five-year horizon, a one-standard-deviation increase in firm innovation raises earnings growth by 2.6 percentage points for workers in the top 5% of the earnings distribution. By contrast, a one-standard-deviation increase in competitor innovation reduces their earnings growth by 4 percentage points. The implied passthrough coefficients—shown in parentheses in the figure—are two to three times larger than those for the average worker. In short, both the gains and the losses from innovation are disproportionately concentrated at the top of the earnings distribution. We next examine whether this heterogeneity also varies systematically across industries.

Industry Heterogeneity

To investigate which workers are more or less exposed to innovation-related income risk, we allow coefficients to vary by industry characteristics. First, we expect the creative destruction channel to be more relevant in firms belonging to industries with high rates of innovation. To illustrate this effect, we re-estimate equations (15) and (16) separately depending on whether the level of industry innovation (the industry analogue of (10)) is above or below the median. As we see in Panel A of Table 1, the magnitudes of the estimated coefficients a and b are significantly larger in industries with high rates of innovative activity. In particular, the response of firm profits, total worker pay and employment to competitor innovation are larger by an order of magnitude in industries in which the rate of innovation is high relative to firms in the other industries. The last two columns show that response of worker earnings for top workers (those in the top quartile) are also significantly larger in industries with high levels of innovation.

Second, we consider two indirect proxies for the labor supply elasticity in our model. First, the within-industry dispersion of earnings: in our quantitative model below, the elasticity of labor supply is endogenously negatively related to the degree of dispersion in workers' productivity across different potential jobs they consider. Second, the average level of earnings: industries with higher average wages are more likely to employ high-skill workers whose labor supply is less elastic. The results in panels B and C of the table are consistent with our intuition. Panel B shows that both own firm and competitor innovation are only significantly related to worker earnings within the subset of industries with high levels of income dispersion. This result obtains despite the fact that own firm innovation is a significant predictor of firm outcomes within the low income dispersion industries. Panel C shows that similar results obtain when comparing industries with higher average wages to those with lower average wages. Together, these results show that innovation shocks transmit unevenly, with the strongest effects in industries where labor supply is likely to be less elastic.

The results so far reveal a striking asymmetry: the empirical passthrough of competitor innovation to worker earnings is roughly twice as large as the passthrough of own-firm innovation—a pattern inconsistent with the symmetric predictions of the simple model in Section 1. We next investigate

the channels behind this asymmetry.

2.3 What Drives the Passthrough of Innovation to Worker Earnings?

The simple model predicts symmetric passthrough because it features only incremental innovation and abstracts from employment transitions. In practice, innovation involves both improving existing products and taking market share from rivals. We now examine three channels that can generate the asymmetry we observe: ex-post heterogeneity in worker outcomes, job destruction, and differences between radical and incremental innovations.

Ex-post Heterogeneity in Worker Outcomes

So far we have focused on average effects. Yet averages mask substantial ex-post heterogeneity across incumbent workers. To uncover this, we estimate equation (16) using the quantile regression approach of Schmidt and Zhu (2016), which characterizes how the entire distribution of earnings growth shifts in response to firm or competitor innovation.

Figure 3 plots the marginal effects—analogous to coefficients a and b—for different quantiles of the earnings growth distribution. The results are stark. Competitor innovation is associated with more negatively skewed subsequent earnings growth for workers. A one-standard-deviation increase in competitor innovation lowers the 5th percentile of earnings growth by 2.2 percentage points, compared to 0.8 points at the median. By contrast, higher innovation by the worker's own firm is associated with more positively skewed earnings growth—a 1.3 percentage point increase in the 95-th percentile compared to a 0.5 percentage point increase in the median earnings growth.

Overall, earnings gains and losses from innovation are concentrated among a small subset of incumbent workers rather than spread evenly. Because gains and losses are concentrated rather than diffuse, the welfare implications for risk-averse workers are substantially worse than mean effects suggest; we quantify this in Section 4.3.

We next examine how these effects vary by worker income. Appendix Figure A.1 shows quantile regression estimates of equation (16) where coefficients a and b vary with the worker's prior income rank. The income gradient in risk is even more pronounced in the tails: top earners face both greater passthrough and greater income risk than the average worker. For example, when competitors innovate, a one-standard-deviation increase reduces the 5th percentile of earnings growth for top workers by 7.1 percentage points—far larger than the 2 point decline at the median. Conversely, when their own firm innovates, top workers see a 3.5 point increase at the 95th percentile, compared to 1.8 points at the median. Thus, high earners not only capture more of the upside from their own firm's innovation but also bear much more downside risk from competitor innovation. A key channel behind these asymmetric outcomes is job loss, which we discuss next.

Job Destruction

A significant component of the passthrough of innovation to incumbent workers' earnings risk arises from the possibility of involuntary job loss, which typically entails large and persistent income costs. To test this channel, we re-estimate equation (16) replacing the dependent variable with an indicator equal to one if the worker experiences an involuntary exit within the next five years. Because the data do not record the reason for separations, we define involuntary exit as a change of employer coinciding with earnings growth below the 20th percentile.

Table 2 reveals an important reason behind the asymmetric passthrough. Competitor innovation raises the likelihood of involuntary exit: a one-standard-deviation increase in competitor innovation increases the probability by 0.81 percentage points. By contrast, a similar increase in own-firm innovation lowers the probability of involuntary exit by 0.33 percentage points. The effects are especially pronounced for high-income workers. A one-standard-deviation increase in competitor innovation raises their likelihood of involuntary exit by 1.87 percentage points, compared to only 0.70 points for lower-paid workers. These results confirm that job destruction is an important mechanism through which innovation shocks translate into earnings risk, particularly for top earners. We next examine whether the type of innovation matters—specifically, whether radical versus incremental innovations have different implications for workers.

Radical vs Incremental Innovations

The discussion above suggests that incremental innovations—improvements to a firm's existing product portfolio—should be associated with higher passthrough to incumbent workers than radical innovations, which are more likely to involve stealing market share from competitors. The logic is that part of a worker's human capital is tied to a specific product; when a firm acquires a new product line, some of the associated gains accrue to newly hired workers rather than incumbents, reducing the benefits passed through to the latter. This mechanism implies an important asymmetry: radical innovations generate larger gains for the firm overall, but a smaller share of those gains reaches incumbent workers. We now examine this prediction in the data.

We classify patents as radical (novel) or incremental using the text-based methodology of Kelly, Papanikolaou, Seru, and Taddy (2021), which measures the pairwise similarity between patents issued to the same firm. A patent i is classified as novel if it is sufficiently distinct from the firm's previous five years of patents:

$$N_i = 1 \Leftrightarrow \max_{j \in P(f)} \rho_{i,j} \leq 0.5, \tag{17}$$

where $\rho_{i,j}$ is the textual similarity between patents i and j, and $P(f)$ is the set of patents filed by firm f in the past five years. Roughly 35% of patents meet this criterion.

At the firm level, we then decompose innovation into radical and incremental components:

$$A^r_{f,t} = \frac{\sum_{j \in P_{f,t}} KPSS_j \, N_j}{B_{ft}}, \qquad A^i_{f,t} = A_{f,t} - A^r_{f,t}. \tag{18}$$

The correlation between A^i_f and A^r_f is approximately 22%.

We then estimate a modified version of equations (15) and (16), decomposing own-firm innovation into radical and incremental parts:

$$g_{i,t:t+h} = a \, A_{I \setminus f,t} + b^r \, A^r_{f,t} + b^i \, A^i_{f,t} + c \, \mathbf{Z}_{ft} + d \, \mathbf{X}_{i,t} + \varepsilon_{i,t}. \tag{19}$$

where i indexes either firms or workers, and the controls are the same as in earlier specifications. For comparability, coefficients b^r and b^i are scaled by the standard deviation of own-firm innovation.

Table 3 presents the results. Panel A shows that radical innovations generate larger increases in firm profits, total labor compensation, and employment than incremental innovations. More importantly, Panel B shows that this pattern does not extend to workers: radical innovations raise profits much more than incremental ones (9.3 vs. 3.9 percentage points), but increase pay for top workers by less (2.3 vs. 2.8 percentage points). The implied passthrough coefficients are 0.17 vs. 0.12 among all workers, and 0.73 vs. 0.25 for workers in the top 5th percentile.

In sum, consistent with the intuition above, incremental innovations pass through more strongly to incumbent workers than radical innovations, even though radical innovations generate larger gains for firms overall. This is in line with the idea that workers' human capital is tied to existing products, so incremental improvements reward incumbents directly, while radical innovations spread gains more broadly to new hires.

Summary

Taken together, these results paint a consistent picture. Firm innovation raises both profits and worker earnings, whereas competitor innovation reduces profits, lowers employment, and depresses worker earnings. The empirical passthrough is *asymmetric*: incumbent workers capture only a modest share of the gains from their own firm's innovation but bear a disproportionately large share of the losses from competitor innovation. This asymmetry is concentrated among high-income workers and in industries where labor supply is relatively inelastic. The channels behind the asymmetry—job destruction, ex-post concentration of gains and losses, and the distinction between radical and incremental innovation—all point to a common mechanism: workers' human capital is tied to specific product lines, so business stealing destroys valuable matches while the gains from expansion accrue partly to new hires.

2.4 Robustness and Additional Results

We next examine the robustness of our findings to alternative measures of competition and earnings, assess the possibility of pre-trends, and compare innovation-based effects to direct measures of

profitability.

Alternative Measures

Competition. Our baseline analysis defines a firm's competitors using 3-digit SIC industries. We now assess whether the results are robust to alternative definitions of competition. Following Hoberg and Phillips (2016), we construct a measure of product-market similarity across firms (TNIC3), which likely provides a more accurate representation of a firm's actual competitors than traditional 3-digit SIC codes. The main limitation of this approach is data availability: TNIC3 is only observable from 1989 onward, which substantially reduces our sample size. As shown in Table A.2, using TNIC3 to define competitor sets yields results that are economically similar to our baseline.

Firm Size. We also re-estimate the analysis using an alternative measure of firm size to scale our innovation measures in (10) and (11)—enterprise value—defined as market capitalization plus book assets minus the value of common equity. Table A.2 shows that results remain quantitatively and economically similar under this specification.

Worker Earnings. Our W-2 earnings measure omits Incentive Stock Options (ISOs), which are untaxed on exercise and therefore absent from W-2 data; 97% of ISO income is concentrated in workers in the top decile of earnings (Eisfeldt, Falato, and Xiaolan, 2022). This omission therefore biases our estimates of passthrough for these top workers towards zero. As a robustness check, we replicate our analysis using adjusted gross income (AGI) and wage and salary income (WSI) from individual tax returns (Form 1040), which include realized gains from stock option exercise. The results are quantitatively similar. Because the 1040-based records are available continuously only after 1999, we use them for robustness rather than as our main analysis. See Panel E of Table A.2 and Appendix B.1 for further details.

Identification

Our identification relies on cross-sectional variation in KPSS patent values as innovation shocks, conditional on an extensive set of firm and worker controls. The key identifying assumption is that, conditional on observables—which include measures of past patenting and R&D investment—the arrival of a highly-valued patent (a patent that elicits a significant increase in firm value) represents an unanticipated shift in a firm's technology. Consistent with this interpretation, the coefficients on firm outcomes in the pre-patent periods ($h = -1$ to $h = -5$) in Figure 1 are neither economically nor statistically significant.

Still, one might worry about unobserved heterogeneity. Firms with persistently higher innovative capacity may attract workers with steeper earnings profiles. Conversely, firms lagging in innovation may be on slower growth paths and employ workers with flatter income trajectories. If our controls for firm R&D, past patenting, and worker earnings history do not fully capture these differences, our estimates could reflect such selection rather than the innovation shock itself.

To address this concern, we ask whether current worker earnings growth predicts future innovation by the firm or its competitors. That is, we re-estimate equation (16) except that we now shift the measures of own-firm and competitor innovation, A_f and $A_{I\setminus f}$, forward,

$$g_{i,t:t+5} = a\, A_{f,t+h} + b\, A_{I\setminus f,t+h} + c\, \mathbf{Z}_{ft+h} + d\, \mathbf{X}_{i,t} + \varepsilon_{i,t}. \tag{20}$$

Our choice of $h = 6 \ldots 10$ avoids any overlap with the five-year period used to measure wage growth. In parallel, we also shift the timing of the firm-level controls $\mathbf{Z}$ to remove any predictable component in innovation. Figure 5 shows no economically or statistically significant relation between wage growth and subsequent innovation. In short, high wage growth in the firm's incumbent workers does not forecast future firm innovation. It is therefore difficult to reconcile our findings with any omitted factor that jointly drives both innovation and worker earnings trajectories.

Alternative Mechanisms

Our interpretation is that the results reflect passthrough of creative destruction in the product market to worker earnings. If this is the case, the impact of innovation on worker earnings should operate through the focal firm's profits. Alternatively, if shifts in industry innovation are driven by changes in product demand in competing firms, then we would expect that increases in competitor innovation are directly associated with lower worker earnings in the focal firm. We cannot include both profits and innovation in the same regression, since profits are themselves an outcome of innovation and thus a bad control (Angrist and Pischke, 2008). Instead, we re-estimate equation (16), replacing the innovation measures $A_{I\setminus f}$ and A_f with profit growth of the competitor and the focal firm, constructed analogously to equations (10) and (11).

Panels A and B of Appendix Figure A.2 show that only profits of the worker's own firm are systematically related to wage growth. Once we control for own-firm profits, changes in competitor profitability have no economically meaningful effect. Replacing profit growth with stock returns yields the same conclusion: if anything, stock returns of competing firms are positively correlated with worker earnings in the focal firm. Such a pattern is consistent with the existence of other common shocks to industries (e.g., demand shocks) which drive both overall industry stock returns and workers' outside options simultaneously. In short, the evidence is consistent with our interpretation that the passthrough of innovation shocks to worker earnings operates through firm profits.

3 The Model

The empirical evidence in Section 2 establishes that the passthrough of creative destruction to worker earnings is asymmetric and concentrated—a pattern the simple model in Section 1 cannot generate. We now build a richer model that reproduces all three features of the data: asymmetric passthrough, concentration of risk among high-income workers, and skewed earnings dynamics.

We extend the simple model along two key dimensions. First, we endogenize labor supply

and introduce rich worker heterogeneity. Firms employ both low-skilled laborers and high-skilled managers, and technology is skill biased do that more productive firms require a higher manager-to-worker ratio. Both types make flexible labor supply decisions along extensive and intensive margins. Following Roy (1951), workers select the best job from a distribution of offers and choose their effort level when employed, generating intensive-margin variation in hours or efficiency units supplied. These features jointly produce realistic heterogeneity across firms and workers.

Second, we incorporate business stealing by allowing firms to invest in two types of innovation. The first is *internal* innovation, which improves the profitability of a firm's existing products, as in Section 1. The second is *external* innovation, undertaken by incumbents and potential entrants, which improves upon a competitor's product and thereby steals its product line, as in Aghion and Howitt (1992). The loss of a product line affects not only firm profitability but also the evolution of workers' human capital: business-stealing events can destroy valuable, high-productivity matches for some workers while creating new, potentially superior, opportunities for others. Our setup enables us to run exact model analogs of the regressions discussed above which facilitate calibration of key parameters of the model via indirect inference, as we discuss in Section 3.4. Section 3.5 provides intuition for how these ingredients jointly generate asymmetric passthrough.

3.1 Static Equilibrium

We begin by outlining the static equilibrium block of the model, in which firms post wages and prices and workers supply labor. Unless noted otherwise, we retain the notation introduced in Section 1. Our discussion therefore highlights only the new elements relative to that baseline.

Production

Firms in each industry compete in producing different varieties as in (1). Following Acemoglu, Gancia, and Zilibotti (2012); Acemoglu et al. (2018), Bertrand competition among potential producers, combined with a small fixed cost of entry, ensures that only the most efficient firm in each variety produces in equilibrium and sets the monopoly price (see Appendix C.1 for details).

A firm with productivity $Q_j(v,t)$ that is active in industry j and variety v produces output using a constant-returns-to-scale technology that combines two types of labor, laborers and managers,

$$Y_j(v,t) = \left((1-B)\ (Q_j(v,t) l_j(v,t))^{\frac{\psi-1}{\psi}} + B\ \left(\bar{Q}_j(t)^{\alpha}\, Q_j(v,t)^{1-\alpha}\ m_j(v,t) \right)^{\frac{\psi-1}{\psi}} \right)^{\frac{\psi}{\psi-1}}, \qquad (21)$$

Here, $l_j(v,t)$ and $m_j(v,t)$ denote efficiency units of laborers and managers, who are complements in production ($\psi \in (0,1)$), so a more productive firm needs both more laborers and more managers. The parameter $\alpha \in [0,1]$ governs skill bias: because managerial productivity is not fully variety-specific, more productive firms demand a disproportionately higher manager-to-laborer ratio, generating strong sorting of high-skilled workers into large firms—similar to Burstein and Vogel (2017). The

leading producer in each variety posts a price $P_j(v,t)$ and piece rates $w_j(v,t)$ and $s_j(v,t)$ per efficiency unit of laborers and managers, expressed relative to aggregate output $Y(t)$; thus, a laborer's per-period earnings equal the number of efficiency units she supplies times $w_j(v,t)\,Y(t)$.

Labor Supply

Workers supply labor along both an extensive margin of job choice and an intensive margin of effort. A worker's labor income equals the product of the number of efficiency units she supplies per unit of effort z, the amount of effort e she chooses, and the posted wage $w_j(v,t)\,Y(t)$ per efficiency unit. Because effort costs are symmetric across all jobs, we can treat the problem sequentially: each worker first chooses the job that offers the highest wage per unit of effort, and then chooses her optimal level of effort.

Workers differ in their productivity across a finite set of potential jobs, and select the job that maximizes earnings, as in the Roy (1951) model. Upon entering the job market, each worker i is randomly assigned a type $\tau(i) \in (\text{Laborer}, \text{Manager})$ with probabilities $(\omega, 1-\omega)$, and observes her match-specific productivity in K jobs (industry–variety pairs)—a tractable stand-in for the information frictions that prevent workers from knowing their productivity in all potential jobs.[1] Let $j(k,i,t)$ and $v(k,i,t)$ denote the industry and variety associated with her k^{th} job, respectively; worker i's match-specific productivity in that job follows a Fréchet distribution,

$$z_{j(k,i,t)}(i, v(k,i,t), t) \sim \text{Fréchet}(T, \theta_{\tau(i)}), \tag{22}$$

where the scale parameter T is the same for laborers and managers, while the dispersion parameters $\theta_{\tau(i)}$ are different. We assume $\theta_l > \theta_m$, so that the productivity of managers is more dispersed than the productivity of workers; this assumption implies that the labor supply of managers is less elastic on the extensive margin than that of laborers. In addition to these K options, each worker also has a fixed outside option (home production) in which she earns $RY(t)$ per unit effort and has productivity $z^*_{i,t} = 1$. We defer the discussion of how worker productivity evolves to Section 3.2.

Each worker chooses the option that maximizes her wage per unit of effort e. We let $z^*_{i,t}$ and $w^*_{i,t}$ denote the productivity and wage, respectively, associated with worker i's optimal choice. Given the Fréchet distribution of productivity z, the probability that a type-$\tau(i)$ worker selects job k from the set of posted wages $\{w_j\}_{j=1}^K$ is

$$P\left[\text{option } k \text{ selected} | \tau(i), \{w_j\}_{j=1}^K\right] = \frac{T_{\tau(i)} w_k^{\theta_{\tau(i)}}}{\sum_{\kappa=1}^K T_{\tau(i)} w_\kappa^{\theta_\tau(i)}} \left(1 - \exp\left(-\left[T_{\tau(i)} \sum_{\kappa \in K} w_\kappa^{\theta_{\tau(i)}}\right] R_{\tau(i)}^{-\theta_{\tau(i)}}\right)\right). \tag{23}$$

Appendix C derives the implied measure of laborers and managers who supply labor to each firm

[1] Our approach is closely related to Bils, Kaymak, and Wu (2025); see also Caunedo et al. (2023); Galle et al. (2023); Hsieh et al. (2019); Burstein et al. (2019); Lagakos and Waugh (2013) for related applications. The K-options assumption is analogous to simultaneous-search models in industrial organization used to capture frictions in consumer choice; see Honka, Hortaçsu, and Wildenbeest (2019) for a survey.

after integrating over the distribution of wages posted by other firms.

We next discuss workers' labor supply decisions along the intensive margin. Workers are hand-to-mouth and have Greenwood, Hercowitz, and Huffman (1988) preferences over consumption c and work effort e,

$$U(c,e) = u\left(c - Y(t)\frac{e^{1+\varphi}}{1+\varphi}\right), \tag{24}$$

where $u(\cdot)$ is a monotonically increasing and concave function satisfying the Inada conditions, and $\varphi > 0$ captures the disutility of effort. The assumption that the disutility of labor is proportional to output $Y(t)$ preserves balanced growth.

Each worker chooses their level of effort to maximize (24). Because each worker consumes her wage earnings in each period, the optimal choice of effort is

$$e^*_{i,t} = (w^*_{i,t} z^*_{i,t})^{\frac{1}{\varphi}}. \tag{25}$$

Workers in better-matched, higher-paying jobs therefore supply more effort. Given (25), the total earnings for worker i in period t can be written as

$$e^*_{i,t}\, w^*_{i,t}\, z^*_{i,t}\, Y_t = (w^*_{i,t}\, z^*_{i,t})^{1+\frac{1}{\varphi}}\, Y_t = c^*_{i,t}. \tag{26}$$

Given the extensive- and intensive-margin decisions described above, the labor supply curves faced by a firm are deterministic functions of its posted wages, derived in Appendix C.1. In contrast to the simple model of Section 1, the quantity of labor supplied per worker is now flexible. By introducing effort on the intensive margin, we add a second channel through which labor supply responds to posted wages. As a result, there are three distinct channels through which labor supply responds to a firm's posted wage. The first is the *intensive margin*: workers in higher-wage jobs supply more effort. The second is *reallocation across employers*: a higher wage draws workers away from competing firms. The third channel operates through *non-employment*: higher wages draw workers out of home production. When a firm needs to attract a very large number of workers, the latter two channels become less important, leaving only the intensive-margin response. Consequently, the labor-supply elasticity faced by an individual firm declines with the wage and asymptotes to $1/\varphi$. This declining elasticity is the source of endogenous monopsony power: more productive firms, which need more labor, face steeper labor supply curves and therefore set larger wage markdowns—a force that generates heterogeneous profit margins and, in turn, heterogeneous passthrough of innovation to worker earnings.

Static Equilibrium

Equilibrium in each period requires that (i) firms choose wages and prices to maximize flow profits, taking other firms' behavior as given, and (ii) firms hold correct beliefs about their competitors' actions. Because all firms are atomistic, no single firm's behavior affects the equilibrium actions of

others.

Panel A of Appendix Figure A.5 shows that more productive firms employ a disproportionately larger share of managers. Although laborers and managers are complements ($\psi < 1$), managerial productivity does not scale one-for-one with firm productivity; as firms expand, they therefore require a higher manager-to-laborer ratio. Panel B shows that the labor supply elasticity for managers is lower than that of laborers on the extensive margin. For a given worker type, more productive firms face a less elastic labor supply curve because they hire more workers in total. In addition, more productive firms employ relatively more managers than laborers. Both of these mechanisms imply that more productive firms face steeper labor supply curves. As a result, we see in Panel C that firm profit margins increase with the firm's productivity advantage relative to competitors. Last, Panel D shows that average wage earnings paid by the firm, for both laborers and managers, increase in firm productivity. Because labor supply is upward sloping and becomes steeper when the firm already employs many workers, more productive firms need to pay a higher wage for both managers and laborers.

In sum, each firm trades off product-market power against the cost of attracting labor from competing employers. Optimal pricing implies that each firm sets a constant markup over marginal cost, but now also exercises monopsony power: because jobs are horizontally differentiated across workers, firms post wages below the marginal product of labor. The effective labor supply elasticity varies across varieties, so more productive firms face steeper labor supply curves, earn higher profit margins, and pay higher wages (see Appendix C.1 for the formal profit equation and a characterization of the static equilibrium properties). This cross-sectional variation in markdowns is the channel through which changes in firm productivity $\hat{q}$ translate into heterogeneous wage responses across the earnings distribution.

3.2 Innovation and Worker Productivity Dynamics

We now describe the dynamic block of the model and delegate all details to Appendix C.2. Appendix Figure A.3 presents the timeline of events within each instant).

Our empirical results exploit time-series variation in the arrival rate of innovation within industries, which indicates that innovation tends to cluster over time. To capture this clustering in a tractable model with a continuum of product lines, we introduce an industry-specific source of risk that produces common shocks to innovation outcomes: an industry-level stochastic clock $\Gamma_j(t)$, a compound Poisson process with non-negative increments:

$$\mathrm{d}\Gamma_j(t) = \mu \, \mathrm{d}N_j(t), \qquad N_j(t+\Delta) - N_j(t) \sim \text{Poisson}\left(\frac{\Delta}{\mu}\right), \tag{27}$$

where $\mathrm{d}N_j(t)$ is a Poisson process with intensity $1/\mu$. Thus, $\Gamma_j(t)$ acts as a subordinator for technological progress: it increases in random increments whose expected size satisfies $E[\mathrm{d}\Gamma_j(t)] = \mathrm{d}t$.

The parameter μ governs the lumpiness of innovation: larger μ produces more volatile industry-level growth, while aggregate quantities evolve deterministically because shocks are independent across the continuum of industries.

Innovation by Incumbent Firms

Similar to Akcigit and Kerr (2018), we allow for two different dimensions of innovation: a firm can innovate within a variety in which it is the leading producer (internal innovation) or in a random variety it does not currently produce (external innovation). To streamline the exposition, we first take firms' innovation investment decisions as given. As in Peters (2020), our assumptions on the cost of innovation will imply that firm investment rates are constant across varieties, which allows us to describe innovation dynamics at the variety level.

Internal innovation. The leading producer in each variety chooses its investment in internal innovation, denoted I, which determines the likelihood of achieving a successful improvement in productivity. An investment level I at time t^- generates a stochastic intensity $I\, d\Gamma_j(t)$ of internal innovation over the next instant. Given the definition of $\Gamma(t)$ in (27), the expected arrival rate of successful innovation equals $I\, dt$.

When a firm succeeds in innovating, its productivity increases by a random number of steps. Each step raises productivity by a factor $\lambda > 1$, and the number of steps reflects the amount of business time that elapses during the innovation burst governed by (27). Thus, conditional on not being displaced by a concurrent external innovation, a firm that invests I experiences a stochastic productivity change in variety v given by

$$\mathrm{d}q_j(v,t) \equiv \mathrm{d}\log Q_j(v,t) = \log\lambda\, N_j^I(v,t) - \delta\,\mathrm{d}t, \tag{28}$$

where δ captures the (smooth) rate of depreciation of the firm's productivity and

$$N_j^I(v,t) \sim \text{Poisson}\big(I_j(v,t^-)\, d\Gamma_j(t)\big)\,, \tag{29}$$

is a variable that counts the number of successful internal innovation steps; it can exceed one because a positive jump in $d\Gamma_j(t)$ represents a burst of business time, allowing multiple innovation steps to occur within an instant.

Therefore, conditional on the incumbent not being displaced by an external innovator, the firm's relative productivity evolves as

$$\mathrm{d}\hat{q}_j(v,t) \equiv \mathrm{d}\log Q_j(v,t) - \mathrm{d}\log \bar{Q}_j(t) = N_j^I(v,t)\log\lambda - \tilde{g}\,\mathrm{d}\Gamma_j(t). \tag{30}$$

That is, a firm's competitive position erodes whenever the rest of its industry innovates faster than it does, and improves only when the firm itself successfully innovates. Internal innovations generate upward jumps proportional to the number of successful steps $N_j^I(v,t)$, while industry-wide innovation bursts raise average productivity $\bar{Q}_j(t)$, creating a downward drift in relative productivity

at rate $\tilde{g}\, d\Gamma_j(t)$.

Appendix Figure A.4 illustrates how the subordinator $\Gamma_j(t)$ maps calendar time into business time, producing clustered innovation bursts. During each burst, the firm's relative productivity $\hat{Q}_j(v,t)$ can shift downward (if industry-average productivity rises) or upward (if the firm itself innovates), generating the lumpy, industry-correlated shocks that our empirical strategy exploits to identify the causal effects of innovation.

Firms choose the level of internal innovation I to equate costs with marginal benefits. The cost of generating new arrivals with intensity I is equal to

$$Y(t)\, c_I\left(I;\hat{q}\right) = \frac{Y(t)}{\varphi_I}\, \Phi\left(\hat{q}\right)\, I^{\zeta}. \tag{31}$$

Following Peters (2020), the function $\Phi\left(\hat{q}\right)$ is chosen so that the equilibrium rate of internal innovation $\bar{I}$ is constant across firms, ensuring tractability by making the innovation problem separable across product lines.

External innovation. The second form of innovation—*external* or *business-stealing* innovation—allows a firm to displace a competitor and take over its product line. This is the source of creative destruction in the model: when a product line is displaced by a competitor, incumbent workers lose their firm-specific match and must draw a new one, generating the concentrated downside risk that drives asymmetric passthrough. External innovation is undirected: the firm draws target varieties at random and improves their productivity by a factor λ. For each product line in which the firm is the leading producer, it can generate business-stealing improvements according to

$$N_j^x(v,t) \sim \text{Poisson}\big(x_j(v,t^-)\, \mathrm{d}\Gamma_j(t)\big)\,, \tag{32}$$

where the firm chooses the intensity $x_j(v,t)$ at cost

$$Y(t)\, c_x(x_j) = \frac{Y(t)}{\varphi_x}\, x_j^{\zeta}. \tag{33}$$

per product line the firm currently operates. These assumptions ensure that the aggregate rate of creative destruction $\bar{x}$ is constant across firms and time. Following Peters (2020), we assume that a small fraction of business-stealing innovations reset the target line's productivity to ensure stationarity of the distribution of relative productivities. This assumption is necessary to ensure that the distribution of relative firm productivities remains stationary, but is also economically sensible: sometimes, firms create new products that completely displace an existing product—as opposed to merely allowing the firm to produce the same product as a rival at a lower marginal cost.[2]

[2]Examples of such displacement include: Xerox 914 replacing carbon-copy and mimeograph workflows; MS word displacing typewriters; and the Apple iPhone displacing the need for standalone MP3 players, GPS units, voice recorders, and compact cameras.

Innovation by New Entrants

A competitive fringe of potential entrants provides a second channel of creative destruction alongside incumbents' external innovation. Each potential entrant pays a flow cost $Y(t)/\varphi_e$ to attempt a business-stealing innovation; free entry ensures that the expected gains from entry equal this cost, pinning down the equilibrium entry rate $\bar{e}$. Each successful entrant arrives with a single product line.

Worker Productivity

Each worker is infinitely lived and begins the model with K potential jobs, with initial match-specific productivity drawn from the Fréchet distribution as in equation (22). Between innovation events, each worker's match quality evolves according to an Ornstein–Uhlenbeck process,

$$\mathrm{d}\, z^*(i,k,t) = -\rho\, z^*(i,k,t)\,\mathrm{d}t + \sqrt{2\rho}\,\mathrm{d}W_{i,k,t}, \tag{34}$$

which gradually pulls match quality toward the mean, capturing depreciation of firm-specific human capital and generating realistic voluntary turnover even absent creative destruction. When an external innovation displaces job k—that is, when the associated product line is lost to creative destruction—the worker receives a new potential job with a fresh productivity draw,

$$z^*(i,k,t) \sim N(0,1), \tag{35}$$

and new industry and variety indices drawn from the product lines that changed their active producer during the innovation burst. Because the worker had previously selected the best-paying job among her K options, this reset typically reduces her best available match quality—the key mechanism generating concentrated downside earnings risk.

3.3 Equilibrium: Dynamics

We begin with a discussion of aggregate dynamics. We then discuss the dynamics of firms, product lines, and worker earnings. Appendix C.2 contains all technical details and a more detailed discussion of the equilibrium.

Aggregate Output

The model features a balanced growth path (BGP) in which the distribution of *relative* productivities $\hat{q}$ is constant across time and across industries. This property follows from the absence of aggregate shocks and the Cobb–Douglas aggregator across varieties in equation (21).

In particular, productivity in industry j evolves according to

$$\mathrm{d}\log \bar{Q}_j(t) = \tilde{g}\,\mathrm{d}\Gamma_j(t) - \delta\,\mathrm{d}t, \tag{36}$$

with growth rate

$$\tilde{g} \equiv (\bar{I} + \bar{x} + \bar{e})(\lambda - 1), \tag{37}$$

which reflects the equilibrium rates of internal innovation $\bar{I}$, external innovation by incumbents $\bar{x}$, and external innovation by entrants $\bar{e}$. Our assumptions on the cost of innovation imply that these equilibrium rates are constant across firms and over time. Aggregate output grows deterministically at approximately $\tilde{g} - \delta$.

Firms and Product Lines

The value of a firm is the sum of the values of all product lines for which it is currently the leading producer:

$$V_{f,t} = Y(t) \sum_{n=1}^{N_{f,t}} V\Big(\hat{q}_{j(f)}(n, f, t)\Big). \tag{38}$$

Here, $V(\hat{q})$ denotes the value of a single product line with relative productivity $\hat{q}$, expressed relative to aggregate output. Because our innovation cost functions make internal and external investment independent across varieties, the firm's dynamic problem is separable across product lines.

The value of a single product line satisfies the Hamilton–Jacobi–Bellman equation

$$\begin{aligned}(r-g)\,V(\hat{q}) \;= \max_{I,x}\Bigg\{ & \pi(\hat{q}) - c_I(I;\hat{q}) - c_x(x) + x\,\mathcal{G}_x\,V \\ & + \frac{1}{\mu}\Bigg[e^{-(I+\bar{x}+\bar{e})\mu} \sum_{k=0}^{\infty} \frac{(I\,\mu)^k}{k!}\; V\Big(\hat{q} + k\log\lambda - \tilde{g}\,\mu\Big) - V(\hat{q})\Bigg]\Bigg\},\end{aligned} \tag{39}$$

The value of a product line incorporates four forces: current flow profits, the option value of improving the line through internal innovation, the possibility of stealing a competitor's line through external innovation (with expected gain $\mathcal{G}_x V$), and the risk of losing the line to creative destruction. The required return on the left-hand side is the interest rate minus the aggregate growth rate, because values are expressed relative to output. Firms equate marginal costs to marginal benefits for both types of innovation, and the entry rate $\bar{e}$ is pinned down by free entry. A detailed economic interpretation of each term and the full derivation of the first-order conditions are provided in Appendix C.2.

Firms choose innovation rates I and x to equate marginal cost to marginal benefit for each type. The equilibrium innovation rates satisfy the first-order conditions

$$c_I'(\bar{I};\hat{q}) = \mathcal{G}_I V \qquad \text{and} \quad c_x'(\bar{x}) = \mathcal{G}_x\,V, \tag{40}$$

That is, firms invest in each type of innovation up to the point where the marginal cost equals the expected gain in product-line value. The operators $\mathcal{G}_I$ and $\mathcal{G}_x$ are defined in equations (C.56) and (C.50) in Appendix C.2. As we discuss in the section 3.2 above, we choose the function $\Phi(\hat{q})$ in the cost of internal innovation investment (31) so that the equilibrium level of internal innovation $\bar{I}$ is independent of $\hat{q}$

Similarly, the equilibrium rate of new firm entry $\bar{e}$ is pinned down by the free entry condition,

which states that entrants cannot make positive profits,

$$\mathcal{G}_e V - \frac{1}{\varphi_e} \leq 0, \tag{41}$$

where $\mathcal{G}_e V$ captures the incremental value of a successful external innovation by an entrant—equation (C.52) in Appendix C.2. Equation (41) holds with equality if entry is positive $\bar{e} > 0$.

Worker Earnings

Workers in the model face two sources of income risk directly tied to firms' innovation decisions. The first operates through wages paid by incumbent firms: when a product line remains active, changes in the firm's relative productivity shift labor demand and therefore wages. Internal innovation raises wages because more productive product lines operate at larger scale and post higher wage rates (Appendix Figure A.5).

The second source of risk is creative destruction. If a product line loses the productivity lead to a competitor, the job disappears and the worker is displaced. Because a worker's human capital z_{ijt} is specific to her product line, displacement forces her to draw a new job with a fresh match-specific productivity draw. Because she had previously selected the highest-paying job among her K options, all remaining alternatives pay less. High-income workers therefore bear the largest losses: they matched into the jobs where their human capital was most valuable.

Appendix Figure A.6 illustrates these two sources of risk by simulating a representative worker's earnings over three years. In the first year, her earnings fluctuate with incremental innovations at her employer. Midway through the second year, her employer is displaced by a competitor; because her previous match was unusually productive, the replacement job offers substantially lower earnings—a sharp, persistent drop. Creative destruction destroys valuable matches, and because high-income workers are those with the best matches, they bear the largest losses.

3.4 Calibration

We first show how to map the key explanatory variables in Section 2—the KPSS value of innovation—into model counterparts, and then describe our parameter choices.

Mapping to Observable Measures of Innovation

To run the same regressions in the model as in the data, we construct model counterparts of the KPSS value of innovation. For firm f, this is the sum of market-value gains across its product lines following an innovation event—which increases through both internal innovation (raising the value of existing lines) and external innovation (adding new lines). At the industry level, the total KPSS value is proportional to $d\Gamma_{j,t}$, the realization of the industry innovation process. This proportionality means that our empirical regressions of worker earnings on competitor innovation identify the causal effect of the industry-level subordinator process described above. The formal expressions for ΔV, $KPSS_{f,t}$, and $KPSS_j$ are provided in Appendix C.2.

Parameter Choice

We calibrate the model in two steps: we first choose a set of parameters based on a priori information, and then calibrate the remaining parameters to match key features of the data.

Parameters calibrated a priori. We first calibrate a subset of parameters to a priori information. The steepness of the cost of internal innovation ζ is set to 2, following Acemoglu et al. (2018) and Peters (2020). We calibrate the degree of innovation uncertainty $\mu = 1/2$ so that, on average, there are two $d\Gamma$ shocks per year, and when they occur innovation progresses as if 6 months of the discrete time process has occurred. We set the number K of job offers that a worker has at any time to ten. See Appendix D for further details.

Parameters calibrated to firm and worker data. We calibrate the remaining 15 parameters to approximately match 36 moments, organized into five categories.

Aggregate targets. We target a growth rate of output equal to 2% along the balanced growth path and an aggregate labor share of 0.6.

Worker-level targets. We match the average level of income risk by targeting the standard deviation of 5-year earnings growth. We target worker mobility by calibrating the model to fit the probability a worker leaves the firm over the next three years, both on average and for each income group, and inequality by targeting each income group's portion of aggregate earnings.

Firm-level regression moments. The next 6 moments come from the first through the third row of Figure 1, and correspond to the response of firm profits, pay per worker, and employment to a KPSS shock by that firm and its competitors at the 5-year horizon. We construct model-implied counterparts to the empirical measure of the value of innovation as described above and run direct counterparts of these regressions in simulated data from the model.

Worker-level regression moments. The next set of 12 moments correspond to the response of individual worker earnings at the 5-year horizon to innovation by the firm and its competitors; we target both the pooled regression coefficients in Figure 2 and the coefficients by worker earnings levels in Figure 4, in practice working with regression coefficients in terms of passthrough from firm profits to worker earnings.

Labor market structure and innovation composition. We target labor supply elasticities by productivity quartile using the estimates in Seegmiller (2024). To pin down the relative contributions of new innovations and existing ones, we calibrate the share in dollars of total innovation stemming from x to 0.36, corresponding to the average dollar share of novel patents as measured by the text similarity measure discussed in Section 2.3. Last, we calibrate the function $\Phi(\hat{q})$ capturing the cost of internal innovation so that the equilibrium amount of internal innovation $\bar{I}$ is independent of firm productivity $\hat{q}$. Appendix Figure A.7 plots the resulting cost function, together with the ergodic distribution of $\hat{q}$.

Parameters. Appendix Table A.4 reports our calibrated parameters. Two features are most noteworthy. First, creative destruction is relatively rare: it occurs at rate $\bar{x} + \bar{e} = 0.084$, implying an average product-line lifespan of about 12 years, with just over $\frac{3}{4}$ of this activity coming from new entrants. Second, managers exhibit far greater skill dispersion than laborers ($\theta_m = 2.13$ vs. $\theta_l = 6.81$), which helps generate heterogeneous passthrough across the income distribution.

Model Fit

Table 4 summarizes the model's fit to the data. The model matches the firm-level responses of profits, employment, and the wage bill to own-firm and competitor innovation (Panel A), and reproduces the heterogeneous passthrough of innovation to worker earnings across income groups, including the steeper gradient and asymmetry for high earners (Panel B; Appendix Figure A.8 shows that the model also matches earnings responses at other horizons). It closely matches the empirical wage-bill distribution and wage inequality proxies (Panel C), quantitatively replicates job-switching rates across income groups (Panel D), and qualitatively matches the relationship between firm productivity and labor supply elasticity from Seegmiller (2024) (Panel E). The model somewhat understates the five-year standard deviation of earnings growth (0.28 vs. 0.57 in the data), while the labor share is quite close to the empirical target (Panel F).

3.5 Model Mechanisms

Two mechanisms jointly generate the asymmetric and concentrated passthrough documented in the data.

Sources of Earnings Risk

Incumbent workers face two sources of income risk directly linked to innovation. The first concerns the passthrough of incremental innovation to earnings and is the focus of the simple model in Section 1. Equation (C.16), which is the generalized version of equation (5) in the simple model, implies that worker earnings respond to changes in a firm's relative productivity $\hat{q}_j(v,t)$ through the firm's labor-demand schedule. Internal innovation raises $\hat{q}_j(v,t)$, increasing labor demand, wages, and the likelihood that the worker remains with her employer. Competitor innovation pushes $\hat{q}_j(v,t)$ downward, lowering labor demand and wages and increasing the probability of job switching. This channel produces broadly symmetric responses to own-firm and competitor shocks.

The asymmetry highlighted in Section 2 between the passthrough of firm and competitor innovation arises from the second source of risk: creative destruction. A worker's human capital z_{ijt} is tied to her product line. When a competitor or entrant overtakes the incumbent, the product line disappears, and the worker must draw a new job with a fresh match-specific productivity draw. Because she previously selected the best-paying option among her K opportunities, all remaining alternatives pay less. Displacement is especially costly for high-income workers, who sort into the

matches where their human capital is most valuable.

The same mechanism also implies that the benefits of own-firm innovation are not fully captured by incumbent workers. Workers benefit when productivity rises in the product lines where they work, but they do not gain from innovations that expand the firm along other lines. For example, when the firm steals a product line from a competitor, much of the gain accrues to newly hired workers rather than incumbents. Panel B of Table 4 reflects this asymmetry: for instance, among the top 5% of earners, passthrough is around 44% for positive shocks but closer to 84% for negative shocks.

Concentration of Earnings Risk

These two sources of risk both fall disproportionately on higher-paid workers for two key reasons. First, higher-income workers are more likely to be managers, who face a fatter-tailed distribution of firm-specific human capital. As Figure A.9 shows, $\theta_m = 2.13$ is roughly one-third of $\theta_l = 6.81$, implying less elastic labor supply and greater wage sensitivity to changes in $\hat{q}$. As a result, their wage earnings are significantly more sensitive to incremental innovation by the firm, or its competitors, than those of laborers.

The second reason is that top earners, either managers or laborers, typically have high idiosyncratic productivity and are employed in high-$\hat{q}$ product lines. These higher-paid workers are particularly vulnerable to job destruction because their current job is typically far better than their next-best option; passthrough rises sharply in the upper tail of the gap distribution (Appendix Figures A.10 and A.11).

4 Model Implications

We now examine the quantitative implications of the model. We first evaluate its performance against empirical moments that are not directly targeted in the calibration; the model's ability to reproduce these patterns provides out-of-sample support for the creative-destruction mechanism. We then ask what the secular decline in radical innovation implies for worker income risk. Finally, we analyze how incomplete markets and uninsurable labor-income risk shape the welfare consequences of innovation policy.

4.1 Non-targeted Moments

We evaluate the model along two dimensions. First, using U.S. administrative data, Guvenen et al. (2021) document that earnings growth is highly skewed—high-income workers experience large, concentrated downside shocks—and that both low- and high-income workers exhibit greater earnings-growth volatility. Second, Seegmiller (2024) shows that firms at the top of the productivity distribution have lower labor shares, reflecting their reliance on workers with relatively inelastic labor supply. We examine the extent to which the model can generate both sets of patterns.

Worker Mobility and Earnings Risk. Figure 6 shows that the model reproduces several key empirical patterns from Guvenen et al. (2021). The model generates excess kurtosis (Panel A) through occasional large displacements, more negative skewness for high-income workers (Panel B) because creative destruction destroys their unusually productive matches, and an increasing earnings-growth variance for top earners (Panel C)—though it fails to match the U-shape at the bottom, in part because we abstract from sources of transitory risk such as discount rate shocks (Meeuwis, Papanikolaou, and Schmidt, 2026). The model also makes empirically plausible predictions for the response of involuntary exit to innovation shocks across income groups (Figure A.12), despite none of these moments being targeted in the calibration.

Labor Share and Productivity. As we saw in Section 3.4, the model also matches the empirical relationship between firm productivity and labor share that arises from wage markdowns. More productive firms face a lower elasticity of labor supply, which leads them to set larger wage markdowns. This implication aligns with the evidence in Seegmiller (2024) that more productive firms have lower labor shares; Appendix Figure A.13 confirms that the model captures this pattern qualitatively across productivity quartiles.

Taken together, these results suggest that the model generates realistic cross-sectional patterns in earnings risk and firm-level labor shares despite not targeting these moments directly.

4.2 Declining Business Dynamism and Income Risk

A well-documented feature of the U.S. economy over recent decades is the decline in business dynamism, particularly in firm entry rates (Haltiwanger, 2012; Akcigit and Ates, 2019). Figure 7 reveals a complementary trend among incumbent firms: the value share of innovation accounted for by radical patents—those that are technologically distinct from the firm's existing base, as defined in equation (18)—has fallen from roughly 40 percent in the 1990s to 20 percent in the 2010s. While the existing literature focuses on the decline in entry, our measure captures a decline in business stealing by incumbents, which corresponds to external innovation x in our model. Together, these trends point to a broad reduction in creative destruction operating through multiple margins.

We next explore what a reduction in the rate of creative destruction implies for firms and workers in our model. We recalibrate the model's share of external innovation—the rate of creative destruction x compared to the total—to match the average radical-innovation share in the 1990s (40 percent) and the 2010s (20 percent). Motivated by the evidence in the literatureon declining entry, we also reduce the rate of entry by the same proportion. We hold all other parameters at their baseline values except for p_τ which we adjust such that the overall fraction of product lines whose productivity is reset to the average remains constant.

In the model, a decrease in the rate of creative destruction lowers the risk of firm displacement and hence lowers earnings risk faced by workers. In addition, however, it increases the dispersion

in productivity across product lines—since it effectively lowers the rate at which the relative productivity of product lines resets to one. These patterns have several implications for firms and workers, which we discuss next.

Labor Share. The observed decrease in the rate of creative destruction implies a decline in the aggregate labor share from 58.8% to 55.3% in our model. Consistent with the data, this decrease in the aggregate labor share is driven by a decline in the labor share among large and highly-productive firms: as the risk of creative destruction falls, firms are more likely to grow larger and more productive. In particular, the log labor share for the firms in the top productivity quartile declines by approximately 0.093 compared to a 0.08 decline in the data (Seegmiller, 2024). By contrast, the log labor share in the firms in the bottom productivity quartile declines by just 0.01 compared to 0.02 in the data. This result concords with evidence that the decline in the labor share is driven by the most productive, "superstar" firms (Autor, Dorn, Katz, Patterson, and Van Reenen, 2020).

Dispersion in Firm Growth Rates. A decline in the risk of creative destruction leads to a decline in the dispersion of firm growth rates. Specifically, our calibrated model implies that the dispersion in 5-year arc employment growth (p90-p10 spread) declined by approximately 0.4 between the 1990s and 2010s, compared to approximately 0.25 in the data (Salgado et al., 2023). In the model, this decrease in dispersion is primarily driven by a decrease in the right tail, a pattern that is consistent with the data (Decker, Haltiwanger, Jarmin, and Miranda, 2016).

Concentration. The increased dispersion in firm productivity leads to increased concentration: in our calibrated model, the revenue share of the top 0.1 percent of firms increases by approximately 15 log points between 1990 and 2010. This increase in concentration is consistent with the data: Kwon, Ma, and Zimmermann (2024) report that the 0.1 percent share of revenue increased by approximately 14 log points over the same period.

Worker Earnings Risk. Moving from the 1990s to the 2010s parameterization reduces the 90 / 10 spread of five-year earnings growth from 0.54 to 0.47—a decline of roughly 7 percentage points. This prediction aligns with the "Great Micro Moderation" documented by Sabelhaus and Song (2010); Salgado et al. (2023), who show that the cross-sectional dispersion of individual earnings growth has declined in an almost secular fashion since 1980, a pattern that is pervasive across demographic groups. In particular, they find a decline in the 90 / 10 spread of men's 5-year earnings growth from around 1.52 to around 1.42 by 2010—a comparable decline in income risk, though the level is quite different between the data and the model. Our model provides a structural mechanism for this fact: the secular decline in creative destruction—operating through reduced radical innovation by incumbents—lowers the rate at which workers are displaced from productive matches, thereby reducing the non-diversifiable component of earnings risk.

Within vs Between-firm Earnings Inequality. Decreasing the rate of creative destruction increases income inequality in the model by increasing heterogeneity in posted wages across firms. This happens because increased dispersion in firm productivity generates more heterogeneity in labor demand across firms, a force which generates between-firm inequality. Our calibration implies that the increase in the between-firm component accounts for more than half (55%) of the increase in variance of worker log earnings levels—compared to 67% in the data (Barth, Bryson, Davis, and Freeman, 2016; Song, Price, Guvenen, Bloom, and Von Wachter, 2019).

In brief, a decline in the rate of creative destruction lowers the degree of uncertainty facing workers and firms (it reduces dispersion in growth rates) yet at the same time it increases inequality among firms and workers (it increases dispersion in levels). This is a partial silver lining of declining business dynamism: while reduced creative destruction may lead to higher income inequality, it also lowers the non-diversifiable earnings risk that workers cannot insure away.

4.3 Welfare

A central implication of standard endogenous growth models (e.g., Romer, 1990) is that innovation generates knowledge spillovers and therefore should be subsidized to raise welfare. That conclusion does not necessarily apply in our setting. Markets are incomplete, and higher innovation increases workers' exposure to labor-income risk. With risk-averse workers and uninsurable earnings fluctuations, the private willingness to pay for additional innovation may be significantly lower than the complete-markets benchmark.

To study this question, we introduce additional structure on worker preferences. Workers have constant-relative-risk-aversion (CRRA) utility,

$$u(C_{it}) = \frac{C_{it}^{1-\gamma}}{1-\gamma}. \tag{42}$$

Workers are hand-to-mouth and consume their wage income. In equilibrium, the consumption index for worker i equals

$$C_{it} = \frac{\varphi}{1+\varphi}(w_{it}z_{it})^{1+1/\varphi}\, Y_t, \tag{43}$$

combining goods consumption and the disutility of effort (see Appendix C.2 for the derivation). We evaluate each worker's willingness to pay—measured as a fraction a_i of consumption over the next $T = 20$ years—for a one-time increase in the innovation rate that raises aggregate growth by $\Delta = 1\%$ for a single year; the formal indifference condition is provided in Appendix D.7. We set $\beta = 0.04$ and $\gamma = 1$ (log utility); given our hand-to-mouth assumption, $\gamma = 1$ is roughly equivalent to $\gamma \approx 3$–4 in a model with a marginal propensity to consume of 0.25–0.30.

We consider three experiments: (i) an increase to all forms of innovation—internal I, external x, and entry e; (ii) an increase in the rate of internal innovation only; and (iii) an increase in the rate of external innovation and entry only. We calibrate these experiments so that aggregate growth rises

by $\Delta = 1\%$ for a single year. We can interpret these experiments as an economy-wide reduction in the cost of innovation, which can occur either as the result of a subsidy or as a result of technological progress (an investment-specific shock). For simplicity, we shift p_τ such that the distribution of wages and firm productivity remains unchanged across the three experiments.

Figure 8 plots workers' willingness to pay for each policy as a function of their initial earnings percentile; the dashed line marks the complete-markets benchmark. Because workers are heterogeneous in productivity and current job offers, there is significant dispersion in willingness to pay. Appendix Figure A.15 shows the results for $\gamma = 2$. We therefore plot the median together with the 25th and 75th percentiles for each earnings bin.

Panel A shows that workers' willingness to pay falls below the complete-markets benchmark, especially for those at the top of the earnings distribution. Workers from the bottom through the 80th percentile are willing to pay 0.87 percent of their consumption for a one-year, 1% increase in the growth rate. For workers above the 80th percentile, willingness to pay declines to 0.23 percent. Averaged across workers, the mean willingness to pay is 0.82 percent of earnings over the next 20 years. Panels B and C reveal that this pattern is driven almost entirely by workers' aversion to creative destruction. Panel B shows that when subsidies apply only to internal innovation, workers' willingness to pay tracks the complete-markets benchmark quite closely—the average is approximately 0.98 percent of earnings. Any dispersion around that benchmark simply reflects differences in match quality with the employer.

Panel C paints a very different picture when we raise the rate of creative destruction, $\bar{x} + \bar{e}$. Workers above the 20th percentile of the earnings distribution exhibit *negative* willingness to pay for additional creative destruction or business dynamism. Those at the 80th percentile would pay roughly 2.9% percent of their consumption *to avoid* subsidizing creative destruction and higher entry, and an average worker in the top decile would pay 12.3% percent. For these workers, creative destruction poses a first-order earnings risk by displacing their firm-specific human capital. By contrast, workers below the 20th percentile gain from increased business dynamism, since displacement raises the odds of drawing a better match than their current one. Averaged across all workers, the mean willingness to pay is -3.7 percent of earnings over 20 years.

Figure 9 extends this analysis by varying the duration of the innovation subsidy from one to twenty years. Panel A shows that when all forms of innovation are subsidized, willingness to pay increases with the length of the subsidy for workers across most of the earnings distribution, but the dispersion across income groups widens substantially. Workers in the bottom 95% of the distribution benefit roughly in proportion to the subsidy's duration, whereas top earners' willingness to pay is non-monotonic: workers in the top 5% of the earnings distribution have negative willingness to pay for subsidies that last between 1 and 10, however magnitudes are attenuated much closer to zero at longer horizons. At the 20 year horizon, the one-time loss from displacement is roughly offset

by the gains from faster year-on-year growth. Panel B performs an analogous calculation but for subsidies which only raise the rate of creative destruction (analogous to Figure 8 Panel C). Here, willingness to pay grows monotonically with the duration of the subsidy, but workers above the median would prefer to avoid such a policy due to the elevated exposure to displacement risk.

In sum, uninsurable labor-income risk sharply reduces the welfare gains from innovation relative to the complete-markets benchmark. The key force is creative destruction: because worker human capital is non-diversifiable, displacement from a high-quality match imposes losses that cannot be insured away. Not all innovation is created equal from a worker welfare perspective—policies that expand the technological frontier without displacing workers are broadly beneficial, while those that rely on business stealing generate concentrated losses for the workers least able to bear them.

These results also carry potential political economy implications. The workers who stand to lose most from subsidies to creative destruction—those with the highest current earnings—are precisely those who command the largest share of current output and, plausibly, the greatest political influence. As emphasized by Acemoglu and Robinson (2000), groups that stand to lose from technological change can use their economic and political resources to block welfare-improving reforms. This asymmetry can generate organized resistance to policies that accelerate business dynamism, even when such policies raise aggregate welfare.

At the same time, these welfare losses are fundamentally *ex-post*—conditional on a worker's current match quality. Behind the veil of ignorance, before workers learn their position in the productivity distribution, willingness to pay for additional innovation coincides with the complete-markets benchmark: higher creative destruction raises the level and compresses the spread of lifetime earnings, so all agents benefit equally in expectation. These ex-ante valuations are reported in the far left columns of Figure 9. The gap between ex-ante and ex-post evaluations echoes Fernandez and Rodrik (1991), who show that reforms which would command majority support after implementation can nonetheless face resistance ex ante when individuals do not know whether they will be among the winners or losers. Our results therefore point to an additional channel—income risk—through which resistance to technological change can arise, especially among higher-paid workers.

5 Conclusion

This paper documents that creative destruction in the product market passes through to worker earnings in an asymmetric and concentrated fashion, and develops a structural model that rationalizes these patterns. The key insight is that worker human capital is tied to individual product lines: incumbent workers bear the full cost of creative destruction but capture only a fraction of its benefits. This mechanism has first-order implications for the welfare evaluation of innovation policy—subsidies that raise internal innovation are broadly beneficial, while those that accelerate creative destruction reduce welfare for most workers. At the same time, from an ex-ante perspective—before workers

learn their position in the match-quality distribution—innovation is unambiguously beneficial, since creative destruction compresses the dispersion of lifetime earnings. The tension between these ex-ante and ex-post evaluations underscores the challenge associated with implementation policies which encourage higher levels of creative destruction and, in turn, concentrated losses on a subset of the population.

Our analysis points to several avenues for future work. First, while the model reproduces many features of the earnings distribution, it understates the unconditional variance of earnings growth. Incorporating additional sources of risk—such as aggregate shocks, on-the-job learning, or human capital accumulation—could help close this gap. Second, our empirical analysis focuses on public firms matched to Compustat, which limits the sample to large, patenting firms. Extending the analysis to private firms or to non-patent-based measures of innovation would broaden the external validity of our findings. Third, the welfare analysis abstracts from the possibility that workers could partially insure against displacement risk through savings, social insurance, or retraining programs. Quantifying how these mechanisms interact with creative destruction is an important question for policy design. Finally, our model treats firm entry and exit as exogenous to labor market conditions. In practice, the availability of skilled labor may itself shape and respond to firms' innovation decisions, creating a feedback loop between labor markets and creative destruction that we leave for future research.

References

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018, November). Innovation, reallocation, and growth. *American Economic Review 108*(11), 3450–91.

Acemoglu, D., G. Gancia, and F. Zilibotti (2012). Competing engines of growth: Innovation and standardization. *Journal of Economic Theory 147*(2), 570–601.e3.

Acemoglu, D. and J. A. Robinson (2000). Political losers as a barrier to economic development. *American Economic Review 90*(2), 126–130.

Aghion, P., U. Akcigit, A. Bergeaud, R. Blundell, and D. Hemous (2018, 06). Innovation and top income inequality. *The Review of Economic Studies 86*(1), 1–45.

Aghion, P., U. Akcigit, A. Deaton, and A. Roulet (2016, December). Creative destruction and subjective well-being. *American Economic Review 106*(12), 3869–97.

Aghion, P., U. Akcigit, A. Hyytinen, and O. Toivanen (2018). On the returns to invention within firms: Evidence from Finland. In *AEA Papers and Proceedings*, Volume 108, pp. 208–12.

Aghion, P., A. Bergeaud, R. Blundell, and R. Griffith (2017). Innovation, firms and wage inequality. Working paper, Harvard University.

Aghion, P. and P. Howitt (1992, March). A model of growth through creative destruction. *Econometrica 60*(2), 323–51.

Akcigit, U. and S. T. Ates (2019). What Happened to U.S. Business Dynamism?

Akcigit, U., J. Grigsby, and T. Nicholas (2017). The rise of American ingenuity: Innovation and inventors of the Golden Age. Technical report, National Bureau of Economic Research.

Akcigit, U. and W. R. Kerr (2018). Growth through heterogeneous innovations. *Journal of Political Economy 126*(4), 1374–1443.

Angrist, J. D. and J.-S. Pischke (2008, December). *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.

Argente, D., M. Lee, and S. Moreira (2024). The life cycle of products: Evidence and implications. *Journal of Political Economy 132*(2), 337–390.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics 135*(2), 645–709.

Autor, D. H., D. Dorn, G. H. Hanson, and J. Song (2014, 09). Trade adjustment: Worker-level evidence *. *The Quarterly Journal of Economics 129*(4), 1799–1860.

Barth, E., A. Bryson, J. C. Davis, and R. Freeman (2016). It's where you work: Increases in the dispersion of earnings across establishments and individuals in the united states. *Journal of Labor Economics 34*(S2), S67–S97.

Bell, A., R. Chetty, X. Jaravel, N. Petkova, and J. Van Reenen (2019). Do tax cuts produce more Einsteins? the impacts of financial incentives versus exposure to innovation on the supply of inventors. *Journal of the European Economic Association*.

Berlingieri, G., M. D. Ridder, D. Lashkari, and D. Rigo (2025, April). Creative destruction through innovation bursts. CEP Discussion Papers dp2095, Centre for Economic Performance, LSE.

Bils, M., B. Kaymak, and K.-J. Wu (2025). Robinson meets roy: Monopsony power and comparative advantage. Technical report, National Bureau of Economic Research.

Braxton, J. C., K. Herkenhoff, C. Huang, M. Nattinger, J. Rothbaum, and L. D. W. Schmidt (2025). Rising risk among the rich: Implications for wealth inequality and interest rates. Working paper.

Burstein, A., E. Morales, and J. Vogel (2019, April). Changes in between-group inequality: Computers, occupations, and international trade. *American Economic Journal: Macroeconomics 11*(2), 348–400.

Burstein, A. and J. Vogel (2017). International trade, technology, and the skill premium. *Journal of Political Economy 125*(5), 1356–1412.

Card, D., A. R. Cardoso, J. Heining, and P. Kline (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics 36*(S1), S13–S70.

Caunedo, J., D. Jaume, and E. Keller (2023, June). Occupational exposure to capital-embodied technical change. *American Economic Review 113*(6), 1642–85.

Chan, M., M. Xu, and S. Salgado (2023). Heterogeneous passthrough from tfp to wages. Working paper, University of Pennsylvania.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2016). Where has all the skewness gone? the decline in high-growth (young) firms in the u.s. *European Economic Review 86*, 4–23. The Economics of Entrepreneurship.

Eisfeldt, A. L., A. Falato, and M. Z. Xiaolan (2022, October). Human Capitalists. In *NBER Macroeconomics Annual 2022, volume 37*, NBER Chapters, pp. 1–61. National Bureau of Economic Research, Inc.

Fagereng, A., L. Guiso, and L. Pistaferri (2018). Portfolio choices, firm shocks, and uninsurable wage risk. *The Review of Economic Studies 85*(1), 437–474.

Fernandez, R. and D. Rodrik (1991). Resistance to reform: Status quo bias in the presence of individual-specific uncertainty. *American Economic Review 81*(5), 1146–1155.

Friedrich, B., L. Laun, C. Meghir, and L. Pistaferri (2019, April). Earnings dynamics and firm-level shocks. Working Paper 25786, National Bureau of Economic Research.

Galle, S., A. Rodríguez-Clare, and M. Yi (2023, 04). Slicing the pie: Quantifying the aggregate and distributional effects of trade. *The Review of Economic Studies 90*(1), 331–375.

Gouin-Bonenfant, E. (2022). Productivity dispersion, between-firm competition, and the labor share. *Econometrica 90*(6), 2755–2793.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, capacity utilization, and the

real business cycle. *American Economic Review 78*(3), 402–17.

Guiso, L., L. Pistaferri, and F. Schivardi (2005). Insurance within the firm. *Journal of Political Economy 113*(5), 1054–1087.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of us workers reveal about lifecycle earnings dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., S. Ozkan, and J. Song (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy 122*(3), 621–660.

Haltiwanger, J. (2012, December). Job Creation and Ind Dynamics in the United States. In *Innovation Policy and the Economy, Volume 12*, NBER Chapters, pp. 17–38. National Bureau of Economic Research, Inc.

Hoberg, G. and G. Phillips (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy 124*(5), 1423–1465.

Honka, E., A. Hortaçsu, and M. Wildenbeest (2019). Empirical search and consideration sets. In *Handbook of the Economics of Marketing*, Volume 1, pp. 193–257. Elsevier.

Hsieh, C.-T., E. Hurst, C. I. Jones, and P. J. Klenow (2019). The allocation of talent and u.s. economic growth. *Econometrica 87*(5), 1439–1474.

Jones, C. I. and J. Kim (2018). A schumpeterian model of top income inequality. *Journal of Political Economy 126*(5), 1785–1826.

Kelly, B., D. Papanikolaou, A. Seru, and M. Taddy (2021, September). Measuring technological innovation over the long run. *American Economic Review: Insights 3*(3), 303–20.

Klette, T. J. and S. Kortum (2004). Innovating firms and aggregate innovation. *Journal of Political Economy 112*(5), 986–1018.

Kline, P., N. Petkova, H. Williams, and O. Zidar (2019). Who profits from patents? rent-sharing at innovative firms. *The Quarterly Journal of Economics 134*(3), 1343–1404.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017). Technological Innovation, Resource Allocation, and Growth. *The Quarterly Journal of Economics 132*(2), 665–712.

Kogan, L., D. Papanikolaou, and N. Stoffman (2020). Left behind: Creative destruction, inequality, and the stock market. *Journal of Political Economy 128*(3), 855–906.

Krebs, T. (2007, June). Job displacement risk and the cost of business cycles. *American Economic Review 97*(3), 664–686.

Kwon, S. Y., Y. Ma, and K. Zimmermann (2024, July). 100 years of rising corporate concentration. *American Economic Review 114*(7), 2111–40.

Lagakos, D. and G. L. Ordoñez (2011). Which workers get insurance within the firm? *Journal of Monetary Economics 58*(6), 632–645.

Lagakos, D. and M. E. Waugh (2013, April). Selection, agriculture, and cross-country productivity differences. *American Economic Review 103*(2), 948–80.

Lamadon, T., M. Mogstad, and B. Setzler (2022, January). Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review 112*(1), 169–212.

Lentz, R. and D. T. Mortensen (2008). An empirical model of growth through product innovation. *Econometrica 76*(6), 1317–1373.

Meeuwis, M., D. Papanikolaou, and L. D. Schmidt (2026, March). Time-varying risk premia, firm insurance, and endogenous labor income risk. Working paper.

Mian, A., L. Straub, and A. Sufi (2021). The saving glut of the rich. Working paper.

Peters, M. (2020). Heterogeneous markups, growth, and endogenous misallocation. *Econometrica 88*(5), 2037–2073.

Romer, P. M. (1990). Endogenous technological change. *Journal of Political Economy 98*(5), S71–102.

Roy, A. D. (1951). Some thoughts on the distribution of earnings. *Oxford Economic Papers 3*(2),

135–146.

Sabelhaus, J. and J. Song (2010). The great moderation in micro labor earnings. *Journal of Monetary Economics 57*(4), 391–403.

Salgado, S., N. Bloom, F. Guvenen, L. Pistaferri, and J. Sabelhaus (2023, September). Evaluating the great micro moderation. Working Paper.

Schmidt, L. D. and Y. Zhu (2016). Quantile spacings: A simple method for the joint estimation of multiple quantiles without crossing.

Seegmiller, B. (2024). Valuing labor market power: the role of productivity advantages. Working paper, Northwestern University.

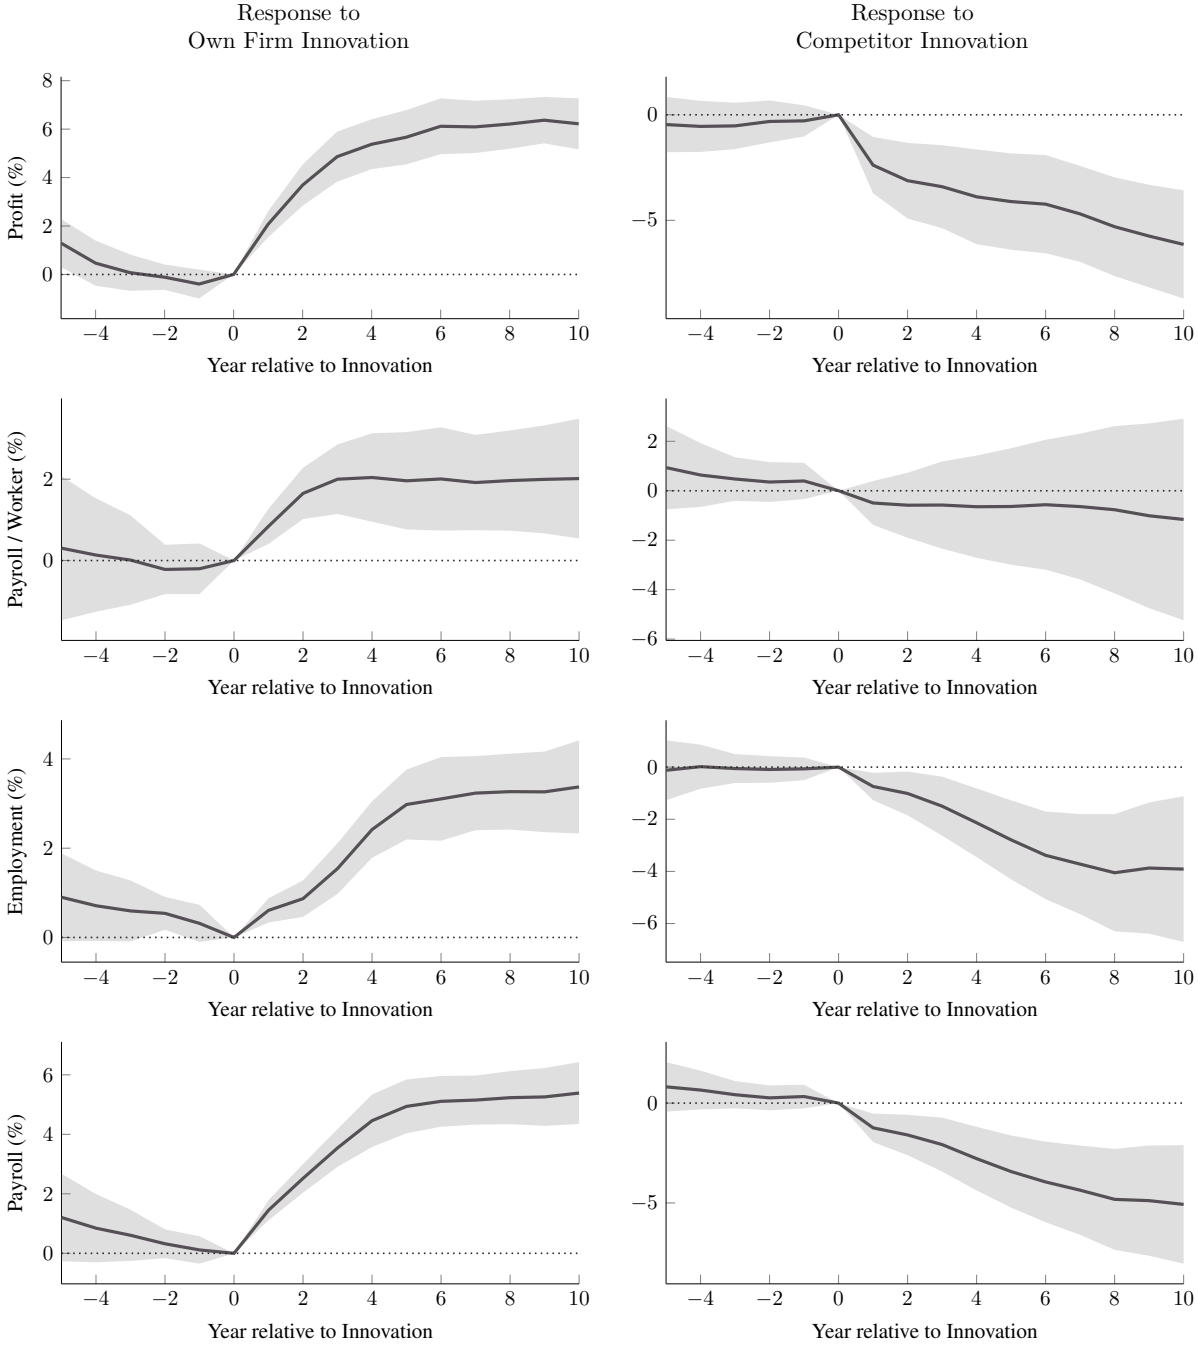
Song, J., D. J. Price, F. Guvenen, N. Bloom, and T. Von Wachter (2019). Firming up inequality. *The Quarterly Journal of Economics 134*(1), 1–50.

Toivanen, O. and L. Väänänen (2012). Returns to inventors. *Review of Economics and Statistics 94*(4), 1173–1190.

van Reenen, J. (1996). The creation and capture of rents: Wages and innovation in a panel of U. K. companies. *The Quarterly Journal of Economics 111*(1), 195–226.

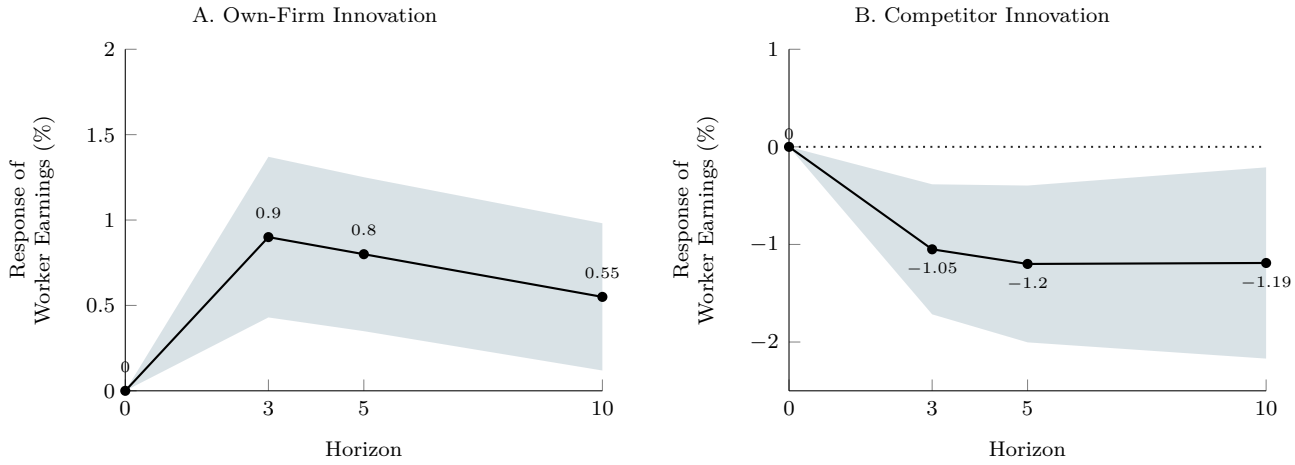
Tables and Figures

Figure 1: Cumulative Changes in Firm Employment, Pay, and Profit Following Innovation



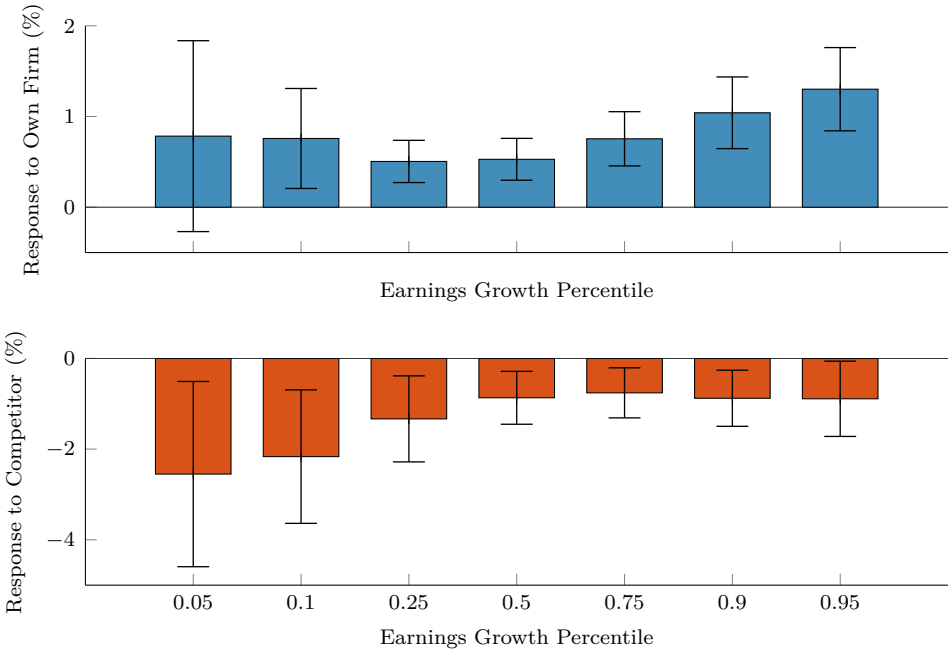
Note: Firm outcomes following a one–standard deviation innovation shock at $t = 0$, shown as cumulative percentage changes. Pay and employment are from the LBD; profit is from Compustat. Regressions control for lagged market capitalization, stock return volatility, assets, patent stock, and industry and year fixed effects. Own-firm innovation is the stock-market surprise to a patent filing; competitor innovation is the sum of such surprises for other firms in the same industry.

Figure 2: Worker Earnings Responses to Own-Firm and Competitor Innovation



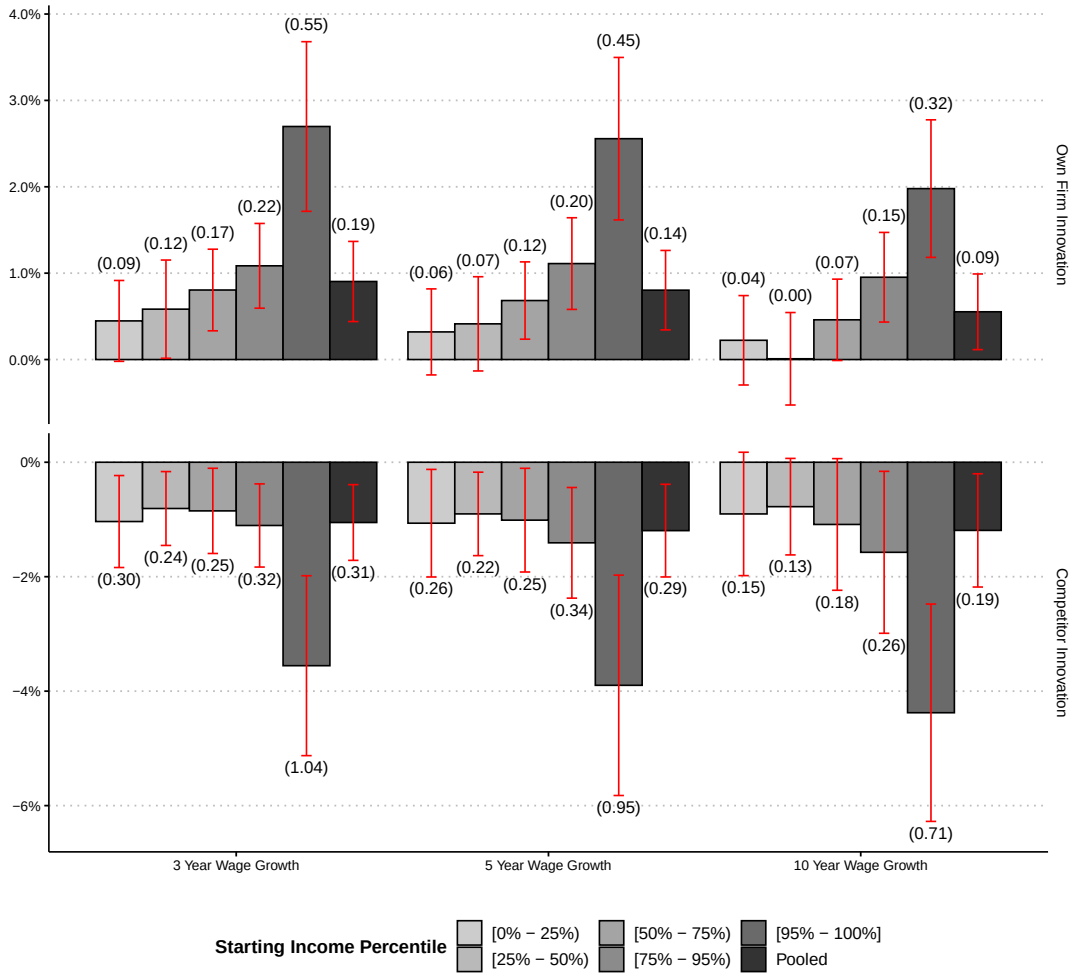
Note: Panels report worker earnings responses at 3-, 5-, and 10-year horizons to own-firm and competitor innovation. Labels above points show estimated effects from the plotted data. Estimation details (controls, weighting, and clustering) follow Section 2.2 and Appendix B.1.

Figure 3: Innovation and the Distribution of Wage Growth: Quantile Regression Estimates



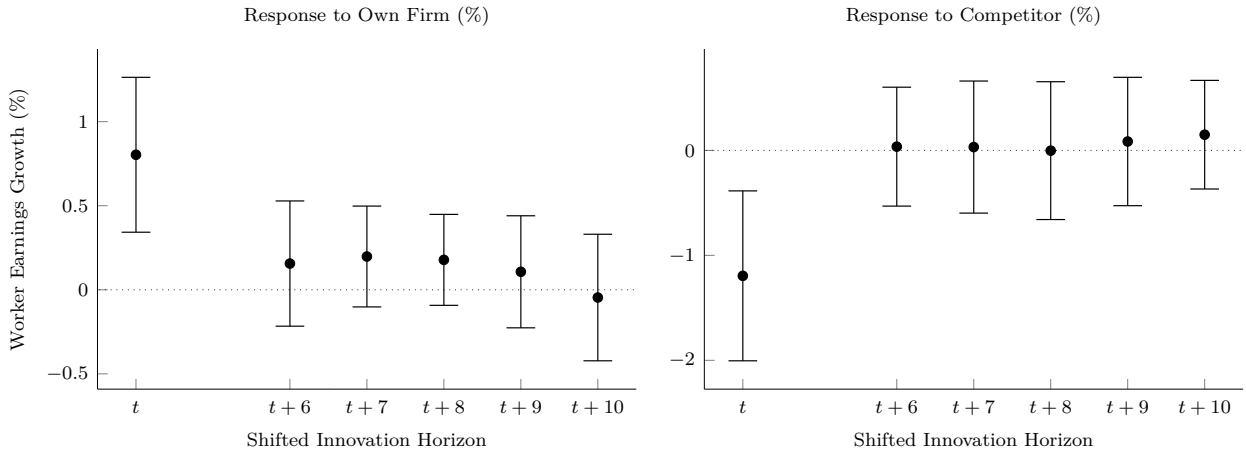
Note: Average marginal effects from quantile regressions (following Schmidt and Zhu 2016) for own-firm and competitor innovation across the earnings-growth distribution. Standard errors use the weighted bootstrap. Variable definitions and full regression specifications are in Section 2.2 and Appendix B.1.

Figure 4: Heterogeneous Responses of Wage Growth to Innovation by Prior Worker Earnings



Note: Heterogeneous worker-earnings responses by prior income group; parenthetical passthrough values scale wage-growth responses by the corresponding 5-year firm-profit responses. See Section 2.2 and Appendix B.1 for full specification details.

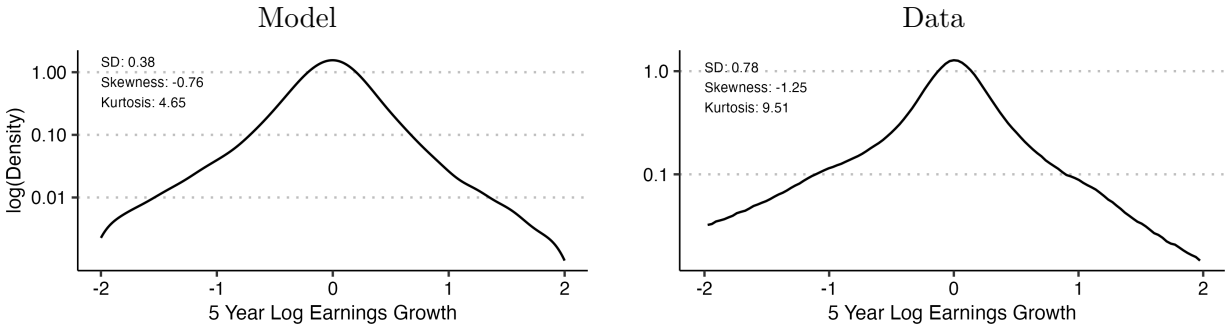
Figure 5: 5-Year Wage Growth and Future Innovation: A Placebo Test



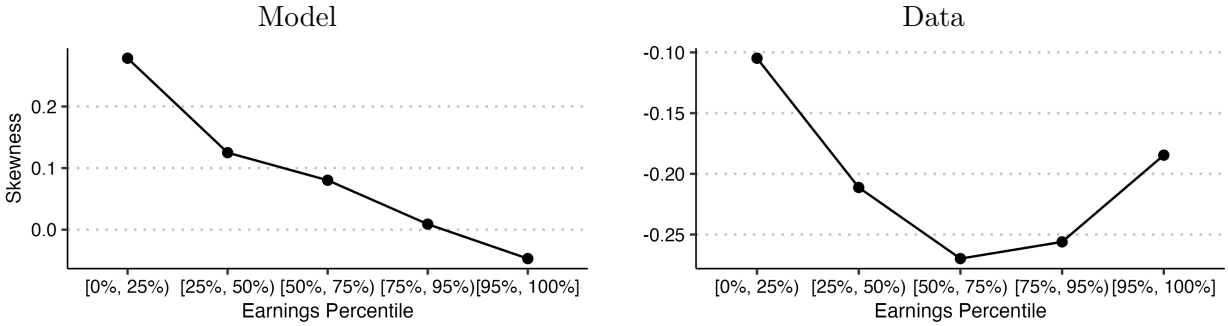
Note: Placebo tests relate past wage growth to future innovation (from lead 6 onward), probing pre-trends and reverse timing. The dependent variable is cumulative worker earnings growth from t to $t + 5$. Specifications follow Section 2.2 and Appendix B.1.

Figure 6: Comparing the Income Process to Guvenen et al. (2021)

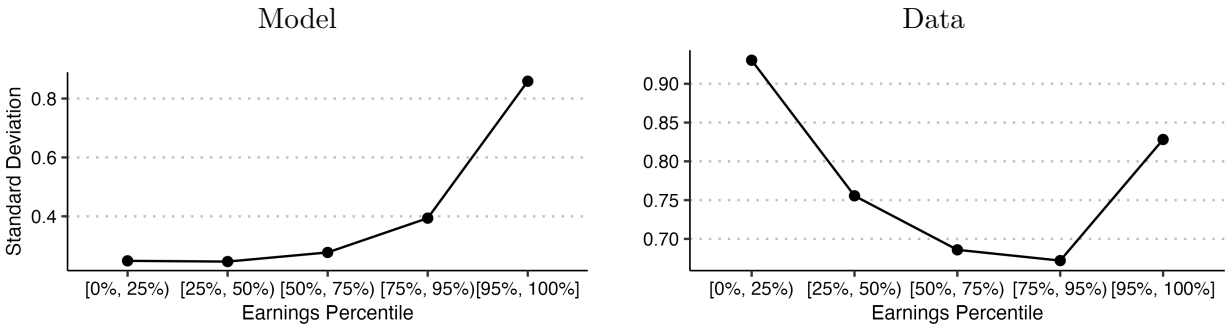
Panel A: Log Density of 5-Year Wage Growth



Panel B: Kelley Skewness of 5-Year Wage Growth by Starting Earnings

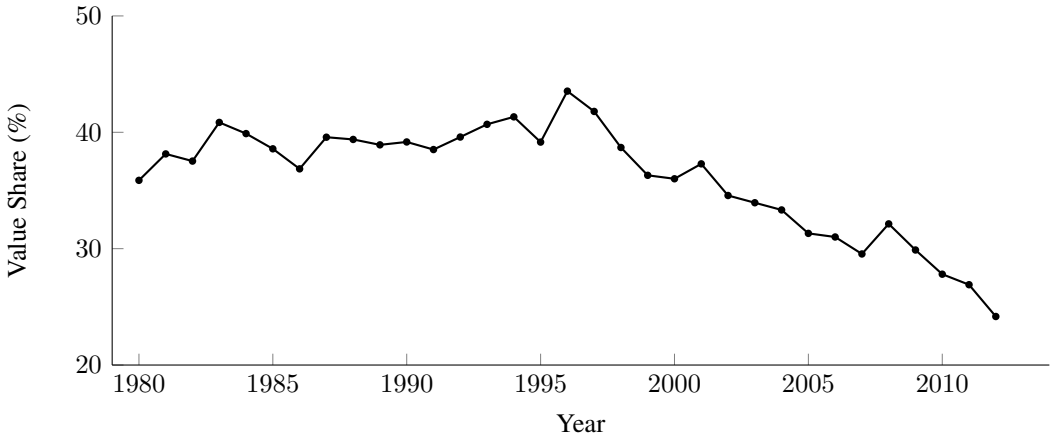


Panel C: Standard Deviation of 5-Year Wage Growth by Starting Earnings



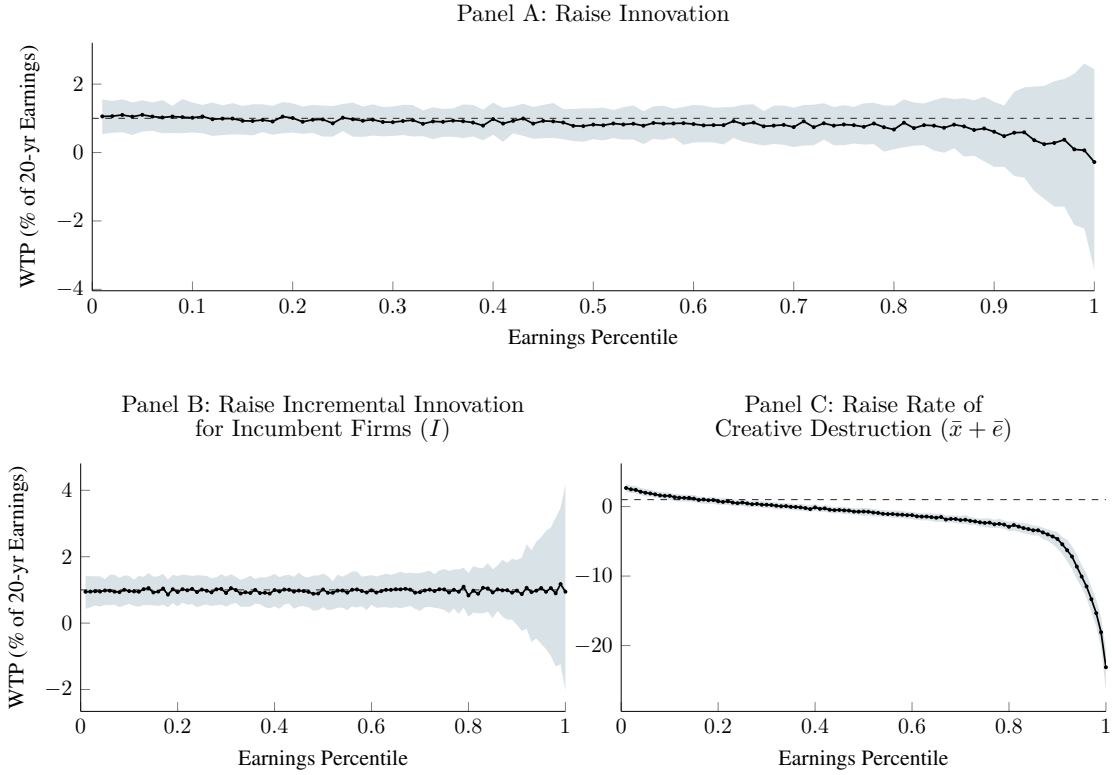
Note: Panel A presents the unconditional log density of 5 year log earnings growth compared with the same numbers as reported by Guvenen et al. (2021), who provide summary statistics by earnings percentile taken from the Social Security Administration data. Panel B shows the Kelly Skewness, computed as the (P90 - P50) - (P50 - P10) for the model vs. the same dataset. Panel A shows the model's implied difference between the 90th and the 10th percentiles of 5-year log wage growth compared to data taken from Guvenen et al. (2021).

Figure 7: Value Share of Radical Innovation Over Time



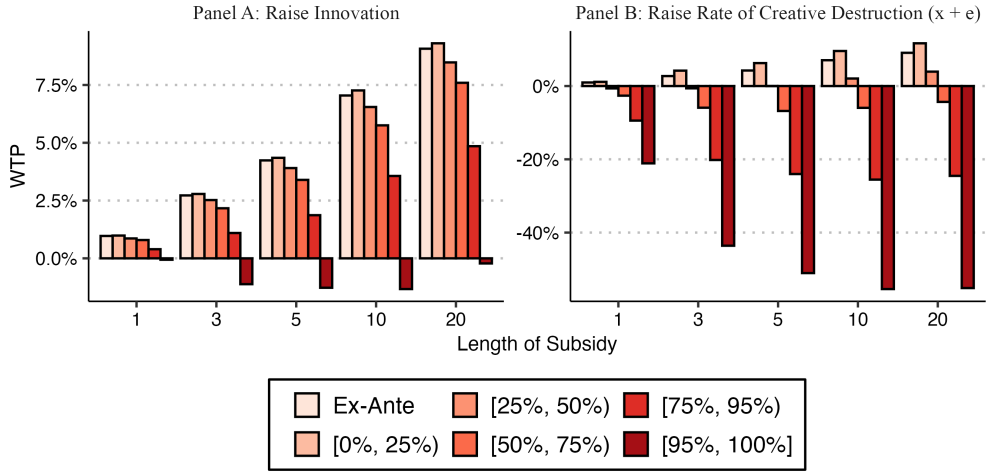
Note: Time series of the relative value (KPSS) of patents that correspond to radical innovation compared to the total value of patents issued in each year. The classification of patents as radical innovation is based on the textual similarity to prior patents issued to the same each firm—see equation (18) in the main text.

Figure 8: Willingness to Pay for a Temporary Increase in Innovation



Note: Willingness-to-pay (WTP) by starting earnings percentile for temporary innovation-policy shocks. Panel A raises both incremental and creative-destruction innovation rates; Panel B raises incremental innovation only; Panel C raises creative destruction only. Lines show medians and shaded bands show interquartile ranges; the dashed line is the complete-markets benchmark. Specification details follow Section 4.3 and Appendix B.1.

Figure 9: Willingness to Pay for Subsidies of Different Lengths



Note: Willingness-to-pay (WTP) by starting earnings percentile for temporary innovation-policy shocks. Panel A raises both incremental and creative-destruction innovation rates; Panel B raises creative destruction only. The bars compare ex-ante WTP under the veil of ignorance as well as average WTP conditional on knowing your own own human capital draws, both pooled and within income percentiles.

Table 1: Firm Outcomes and Wage Growth for Industry Subgroups

		5 Year Firm Outcomes			5 Year Wage Growth	
		Profit	Pay	Employ.	P50–75	P95–100
Panel A. Industry Innovation Rate						
Own Firm	High	8.790	7.310	4.430	1.230	3.780
		(0.775)	(0.652)	(0.739)	(0.304)	(0.732)
	Low	2.510	2.180	1.480	-0.099	0.235
		(0.714)	(0.580)	(0.612)	(0.124)	(0.267)
Competitor	High	-4.560	-3.970	-2.470	-0.823	-4.270
		(1.570)	(1.190)	(1.040)	(0.621)	(1.300)
	Low	-0.183	-0.200	-0.079	-0.375	-0.438
		(0.432)	(0.439)	(0.399)	(0.145)	(0.203)
Panel B. Industry Income Inequality						
Own Firm	High	7.190	6.030	3.640	0.869	2.930
		(0.613)	(0.546)	(0.613)	(0.261)	(0.588)
	Low	3.800	2.240	1.670	0.043	0.271
		(0.797)	(0.580)	(0.625)	(0.140)	(0.230)
Competitor	High	-4.310	-4.000	-2.670	-1.030	-4.440
		(1.340)	(1.040)	(0.900)	(0.585)	(1.220)
	Low	0.380	-0.137	-0.265	-0.026	-0.062
		(0.520)	(0.628)	(0.641)	(0.230)	(0.344)
Panel D. Average Industry Wage						
Own Firm	High	7.540	6.380	3.810	1.030	3.440
		(0.652)	(0.577)	(0.624)	(0.271)	(0.579)
	Low	2.980	1.760	1.750	0.269	-0.646
		(0.707)	(0.512)	(0.725)	(0.180)	(0.492)
Competitor	High	-4.240	-4.010	-2.700	-0.784	-4.810
		(1.410)	(1.070)	(0.942)	(0.527)	(1.040)
	Low	-0.656	-0.764	-0.711	0.448	-0.605
		(0.456)	(0.531)	(0.580)	(0.447)	(0.849)

Note: This table contains responses to own firm and competitor innovation, where industries are grouped into high or low relative to the median industry. All coefficients are multiplied by 100. Industry concentration is based on HHI of Compustat sales. Industry income inequality is based on the cross-sectional variance within an industry in a given year. Industry innovation rate is measured by the rate of an industry's patenting activity. Industry wages are measured using the CPS-DER linked sample.

Table 2: Effects of Innovation on Likelihood of Involuntary Exit

Pr(Involuntary Exit)	Pooled	By earnings group				
		P0–25	P25–50	P50–75	P75–95	P95–100
Own Firm Innovation	-0.33%	-0.17%	-0.22%	-0.31%	-0.42%	-0.79%
	(0.12%)	(0.14%)	(0.18%)	(0.12%)	(0.09%)	(0.22%)
Competitor Innovation	0.81%	0.70%	0.74%	0.75%	0.89%	1.87%
	(0.29%)	(0.32%)	(0.33%)	(0.33%)	(0.30%)	(0.41%)

Note: Involuntary exit is defined as leaving a firm and experiencing 5 year wage growth in the bottom 20% of the sample for that given year. Columns show probability of involuntary exit in response to a one standard deviation change in innovation based on the estimates from a linear probability model. Controls are identical to the other worker-level regressions in the text, and standard errors (shown in parentheses) are clustered by industry and year.

Table 3: Firm vs Worker Responses to Radical and Incremental Innovation

	A. Firm Outcomes			B. Worker Wage Growth					
	Profit	Total Payroll	Employment	Pooled	[0%, 25%)	[25%, 50%)	[50%, 75%)	[75%, 95%)	[95%, 100%]
Own Firm, Radical (novel)	9.25%	7.70%	5.59%	1.09%	1.29%	0.12%	0.96%	1.50%	2.31%
	(0.99%)	(0.65%)	(0.71%)	(0.52%)	(0.74%)	(0.56%)	(0.68%)	(0.77%)	(1.17%)
				[11.79]	[13.99]	[1.29]	[10.43]	[16.18]	[24.95]
Own Firm, Incremental (non-novel)	3.88%	3.36%	1.21%	0.66%	-0.26%	0.73%	0.56%	0.81%	2.82%
	(1.01%)	(0.48%)	(0.54%)	(0.32%)	(0.50%)	(0.46%)	(0.46%)	(0.39%)	(0.66%)
				[17.08]	[-6.72]	[18.77]	[14.42]	[21.00]	[72.59]
Competitor	-4.24%	-3.98%	-3.42%	-1.20%	-1.07%	-0.89%	-1.01%	-1.42%	-3.91%
	(1.20%)	(1.04%)	(0.87%)	(0.41%)	(0.48%)	(0.37%)	(0.46%)	(0.49%)	(1.00%)
				[28.19]	[25.12]	[21.10]	[23.83]	[33.50]	[92.12]

Note: Novel patents are measured based on textual similarity to prior patents. This table compares firm and worker outcomes on a 5 year horizon for novel and non-novel patents. Standard errors are in parentheses, and the implied passthrough from profits to wages is in brackets. Passthrough (wage growth relative to firm profit responses) is substantially higher for non-novel innovation, especially for higher income individuals, since wage growth is larger but profit responses are smaller for non-novel patents relative to novel ones.

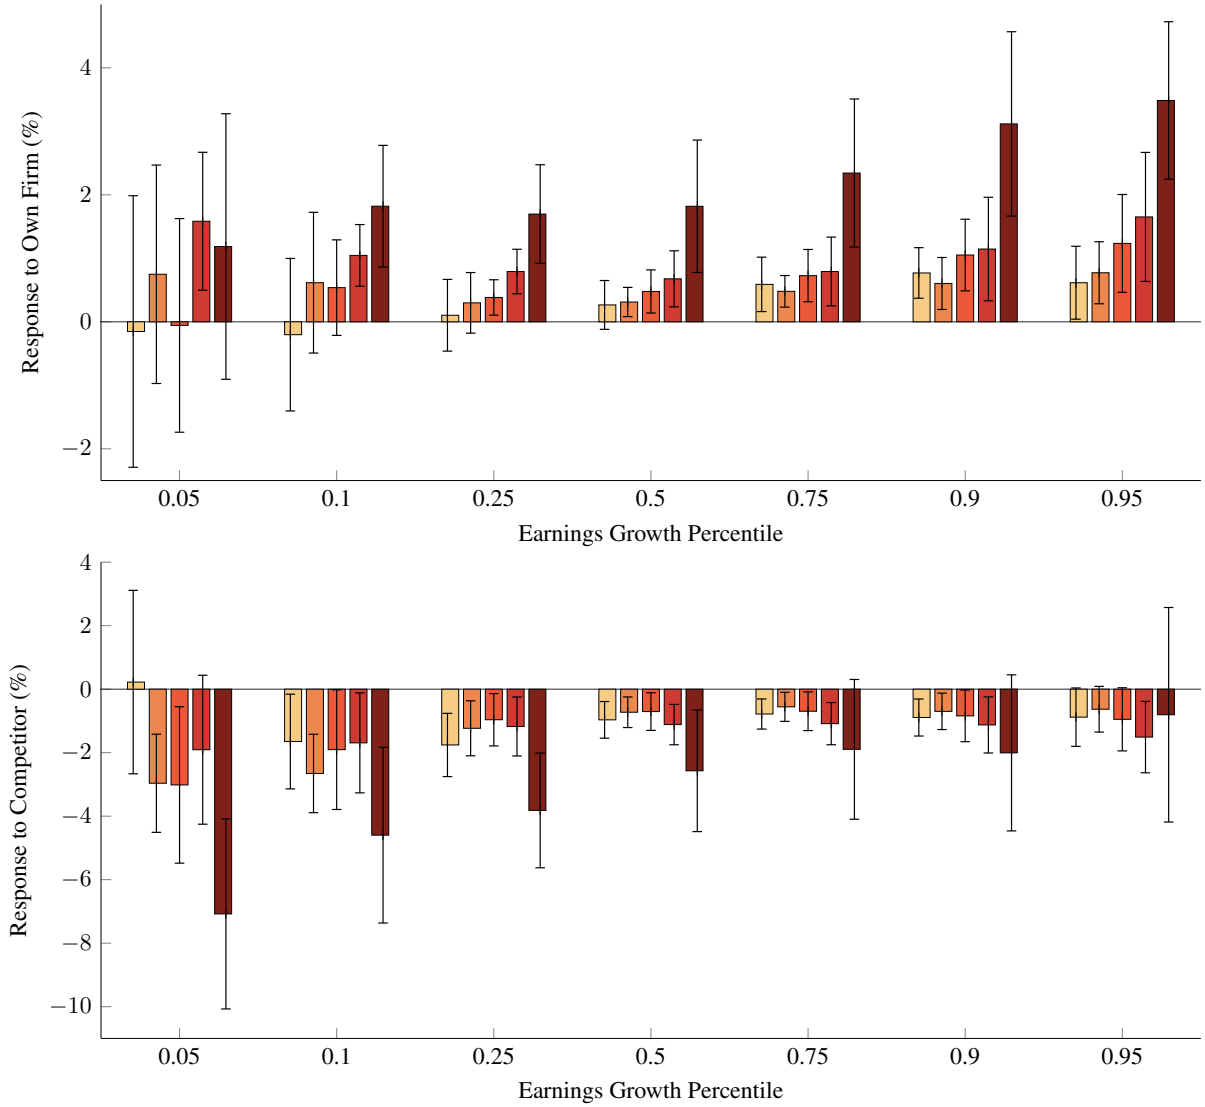
Table 4: Targeted Moments in Model Calibration

Moment	Model	Data	Source
A. Firm Responses to Innovation (5-years, %)			
Profitability			Figure 1
Response to Firm Innovation	9.39	5.66	
Response to Competitor Innovation	-3.29	-4.11	
Employment			Figure 1
Response to Firm Innovation	7.54	2.98	
Response to Competitor Innovation	-2.65	-2.79	
Payroll			Figure 1
Response to Firm Innovation	8.11	4.94	
Response to Competitor Innovation	-2.85	-3.43	
B. Passthrough to Worker Earnings (Worker Earnings / Firm Profits, 5-years)			
Pooled			Figures 1 and 2
Response to Firm Innovation	0.11	0.14	
Response to Competitor Innovation	0.23	0.29	
[0–p25]			Figures 1 and 4
Response to Firm Innovation	0.06	0.06	
Response to Competitor Innovation	0.07	0.26	
[p25–p50]			Figures 1 and 4
Response to Firm Innovation	0.09	0.07	
Response to Competitor Innovation	0.16	0.22	
[p50–p75]			Figures 1 and 4
Response to Firm Innovation	0.11	0.12	
Response to Competitor Innovation	0.22	0.25	
[p75–p95]			Figures 1 and 4
Response to Firm Innovation	0.19	0.20	
Response to Competitor Innovation	0.43	0.34	
[p95–p100]			Figures 1 and 4
Response to Firm Innovation	0.43	0.45	
Response to Competitor Innovation	1.00	0.95	
C. Earnings Inequality (Earnings Share, %)			
[0–p25]	16.4	11.0	Detailed Earnings Record
[p25–p50]	19.4	17.0	
[p50–p75]	22.8	24.0	
[p75–p95]	23.3	29.0	
[p95–p100]	18.1	18.0	
D. Job Switching Probabilities (3-years)			
Pooled	51.8	41.0	Detailed Earnings Record
[0–p25]	65.1	50.3	
[p25–p50]	55.4	40.8	
[p50–p75]	47.9	37.1	
[p75–p95]	38.5	35.7	
[p95–p100]	31.5	37.6	
E. Labor Supply Elasticities by Firm Productivity Quartile			
[0–p25]	3.62	4.32	Seegmiller (2024)
[p25–p50]	2.86	2.71	
[p50–p75]	2.39	2.34	
[p75–p100]	1.68	1.17	
F. Other			
Labor Share	0.58	0.6	
Standard Deviation of 5-Year Wage Growth (%)	28	57	Detailed Earnings Record
KPSS Value share of Innovation coming from x / Radical Innovation (%)	34.6	34.9	Kogan et al. (2017) and equation (18)
90% - 10% spread in log revenue per product line	3.98	3.90	**?**

Note: Moments sourced from the Detailed Earnings Record refer to calculations by the authors done with the regression sample used in this paper, with worker earnings drawn from the Detailed Earnings Record as described in Section 2.1. Moments based on data from Kogan et al. (2017) are using data on the dollar value of patents from the aforementioned paper.

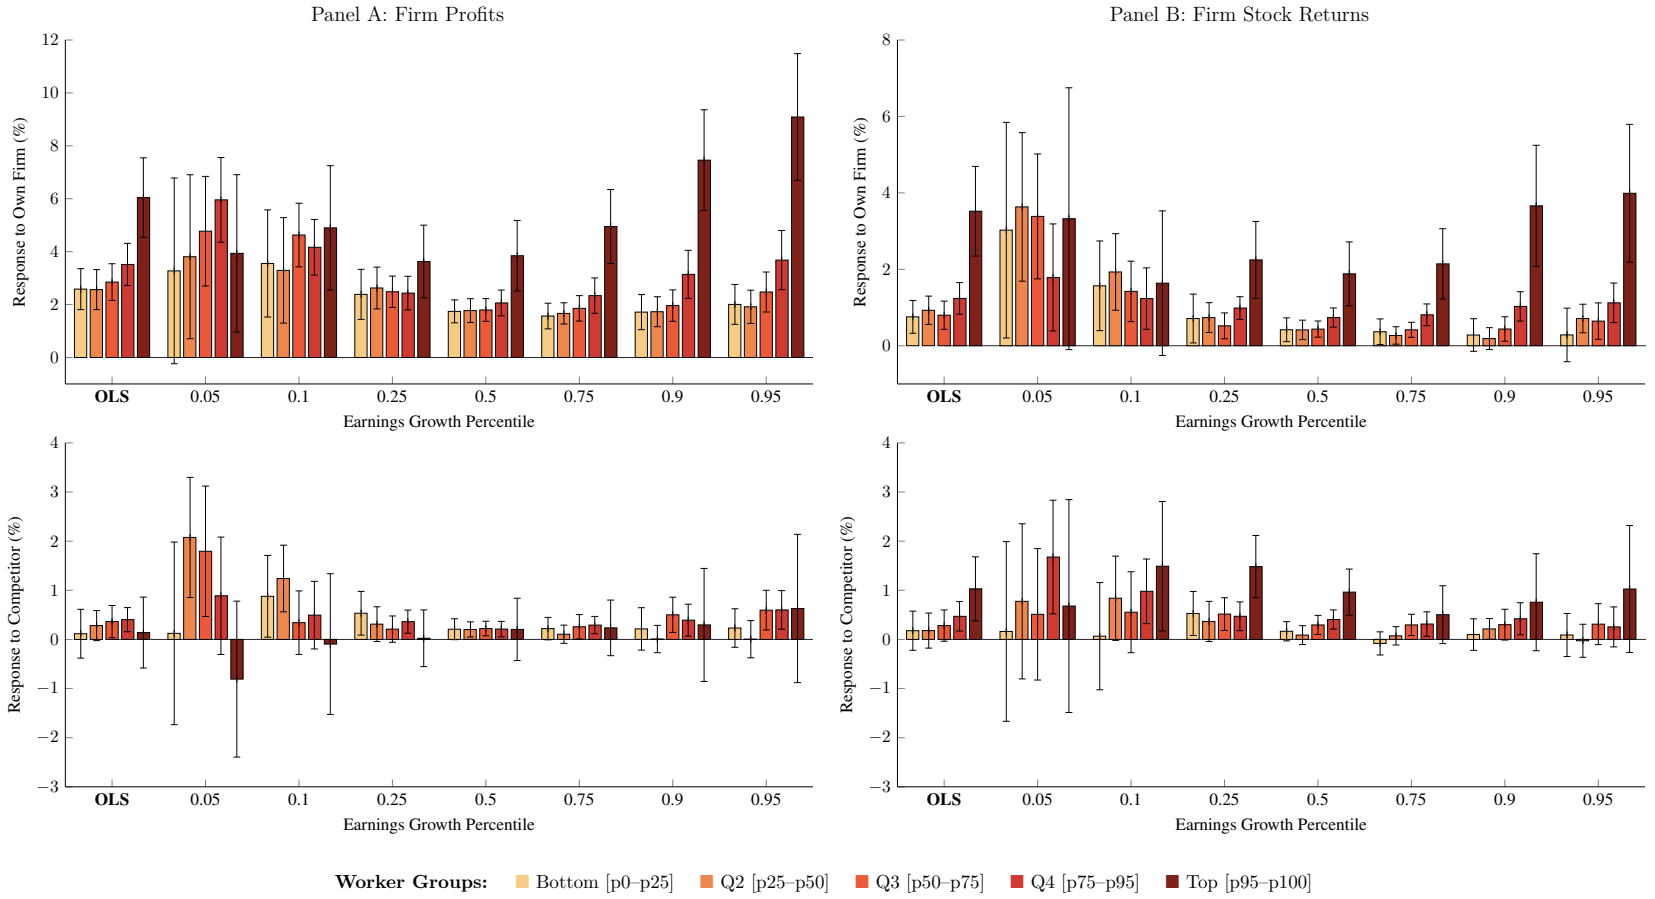
A Appendix Tables and Figures

Figure A.1: Innovation and the Wage Growth Distribution by Income Group: Quantile Regression Estimates



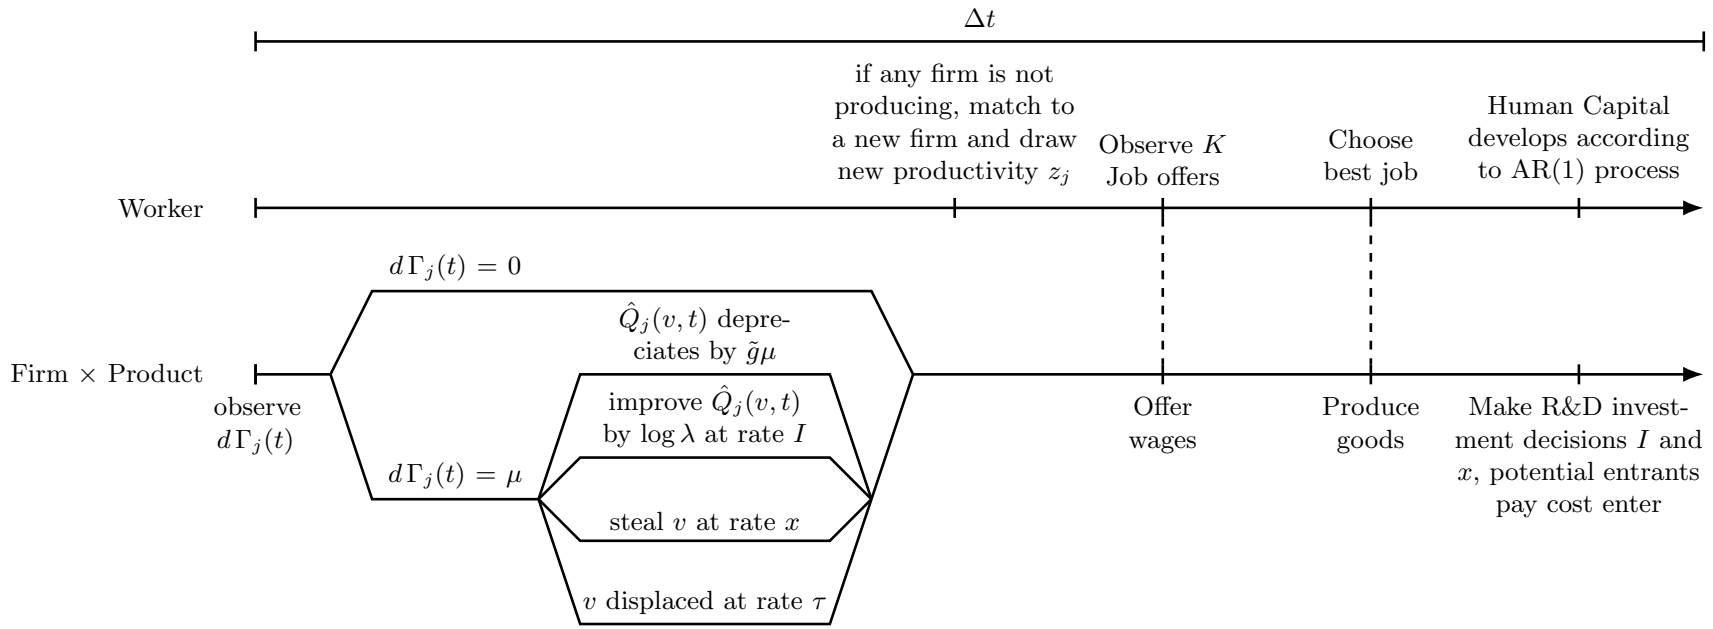
Note: The figures report average marginal effects from quantile regression, estimated using the method of Schmidt and Zhu (2016), which guarantees that fitted quantiles do not cross. Standard errors are estimated using the weighted bootstrap. Coefficients are interacted with income bins, showing substantially larger responses for high income individuals on the up- and down-sides, with the left tail growing substantially following competitor innovation. Controls include lagged firm market capitalization, volatility of stock return, assets, and the lagged, depreciated value of a firm's patent stock at the firm level, lifecycle effects of wage paths approximated by Chebyshev polynomials, lagged wage growth, and industry cross income bin and year cross income bin fixed effects. Income bins are indicator variables for whether a person's income is within one of the following quantiles within a given industry and year: $[0\%, 25\%)$, $[25\%, 50\%)$, $[50\%, 75\%)$, $[75\%, 95\%)$, and $[95\%, 100\%]$. Own-firm innovation is measured by the surprise component of the stock market response to a patent filing, and competitor innovation is measured by the sum of these responses to a patent approval for an entire industry, leaving out the "own firm."

Figure A.2: Passthrough of Firm Profits and Stock Returns to Worker Wage Growth by Income Group



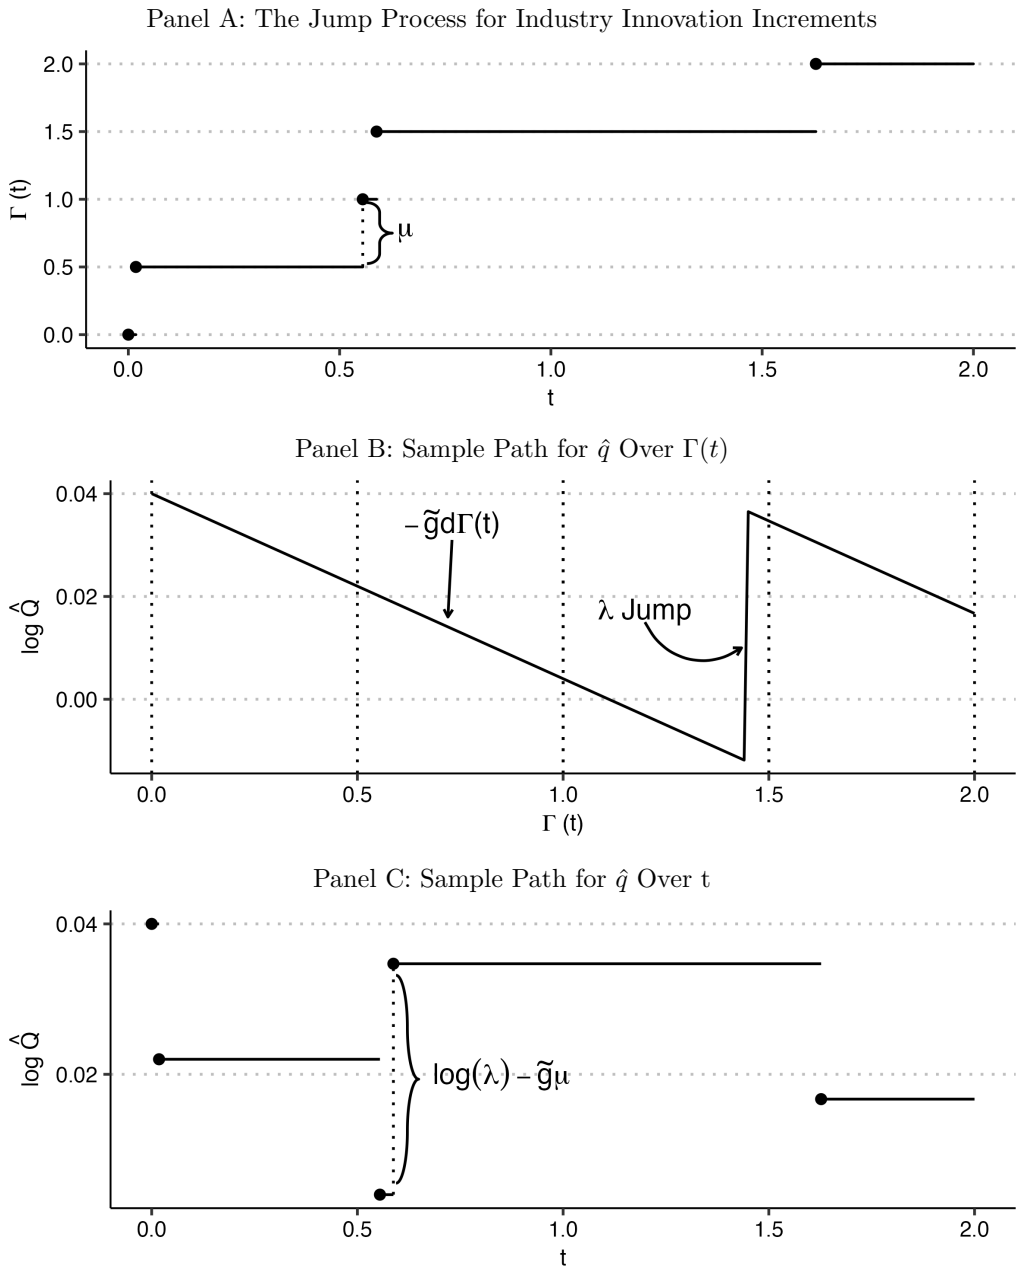
Note: The figure shows responses of five year wage earnings growth to a shock to own firm and competitor profits (Panel A) and stock returns (Panel B). The figure shows both OLS responses, and the estimates from quantile regressions. Controls include lagged firm market capitalization, volatility of stock return, assets, and the lagged, depreciated value of a firm's patent stock at the firm level, lifecycle effects of wage paths approximated by Chebyshev polynomials, lagged wage growth, and industry cross income bin and year cross income bin fixed effects. Income bins are indicator variables for whether a person's income is within one of the following quantiles within a given industry and year: $[0\%, 25\%)$, $[25\%, 50\%)$, $[50\%, 75\%)$, $[75\%, 95\%)$, and $[95\%, 100\%]$. As before, the definition of competitor is the NAICS4 industry of the firm aggregate, less the "own firm."

Figure A.3: Timeline of worker and firm actions within a period.



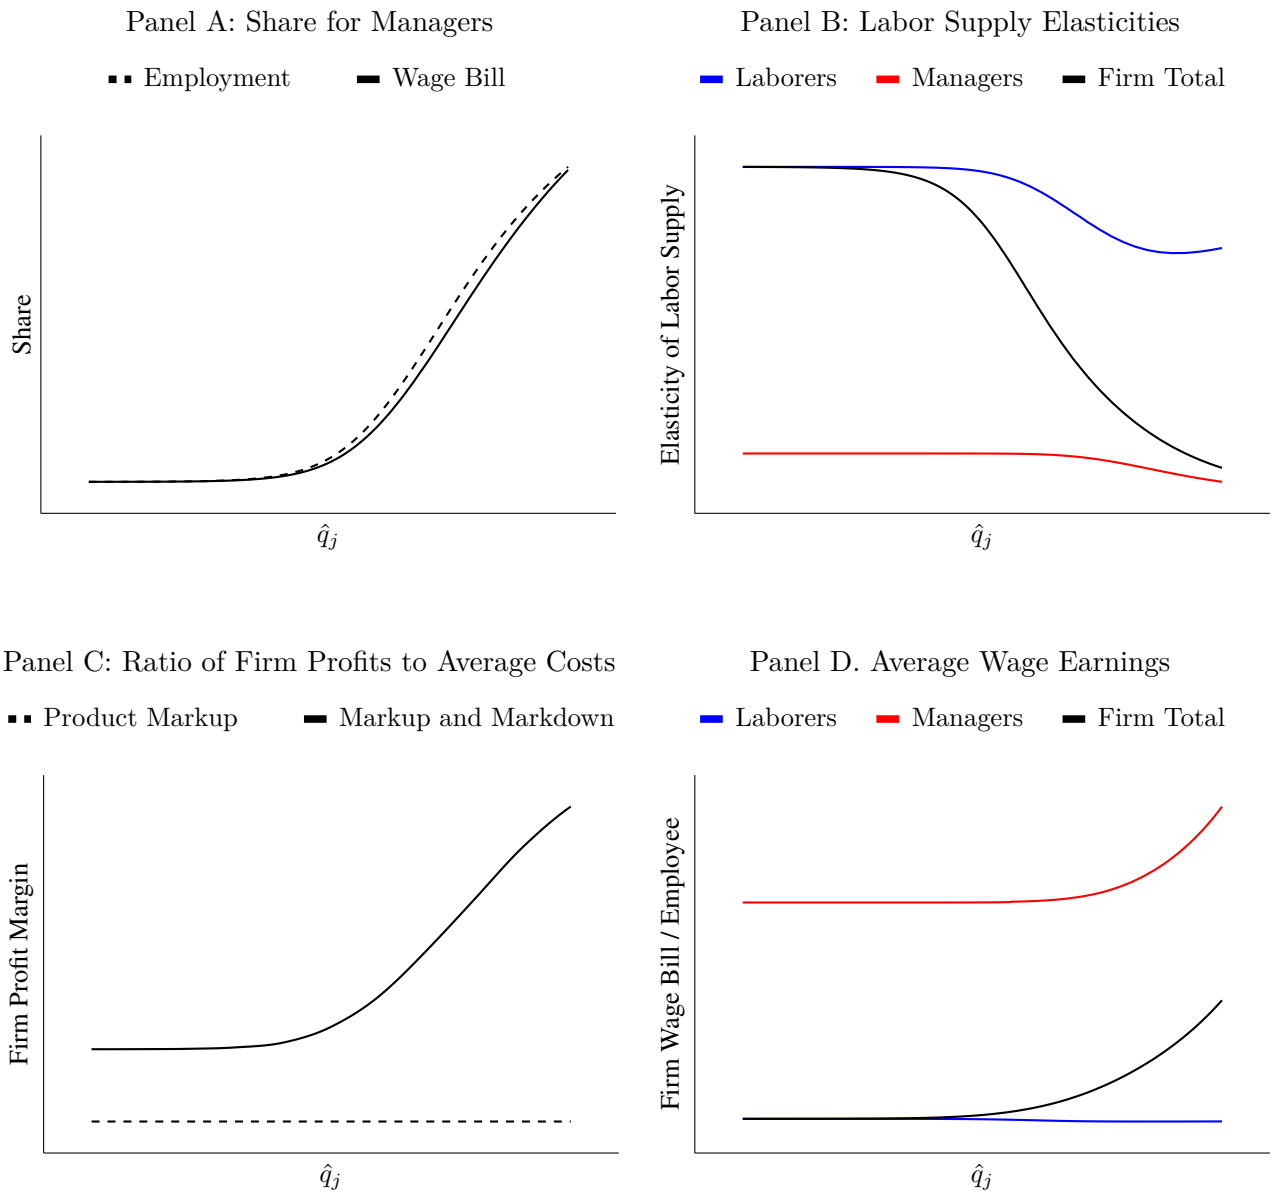
Note: The above describes the actions taken by the workers and firms over a small time interval Δt. The firm first observes if an industry shock $d\Gamma_j(t)$ occurs. For a time shock μ, the relative efficiency of the firm $\hat{Q}_j(v,t)$ depreciates by $\tilde{g}\mu$ as the rest of the industry improves. The firm efficiency improves according to a Poisson process with rate I. Similarly, the firm is displaced from varieties of production at rate $\tau = x + \varepsilon$ but steals varieties from other firms with rate x. If there is no time shock, then the varieties and efficiencies of all firms remains constant. Following these developments, the firm offers wages to the distribution of workers. Each worker observes K job offers and picks the best one. If any of the K firms is entirely displaced and no longer produces anything, the worker is matched to a new firm. Once the agents picks their best job, the firm produces goods. At the end of the time increment, the human capital of the worker for each of the K jobs develops along an AR(1) process and the firm decides how much to invest in R&D for the next period. A mass potential entrants with measure ε decide whether to pay an entry cost to enter the industry.

Figure A.4: Time for the Innovation Process vs. Time for the Observer



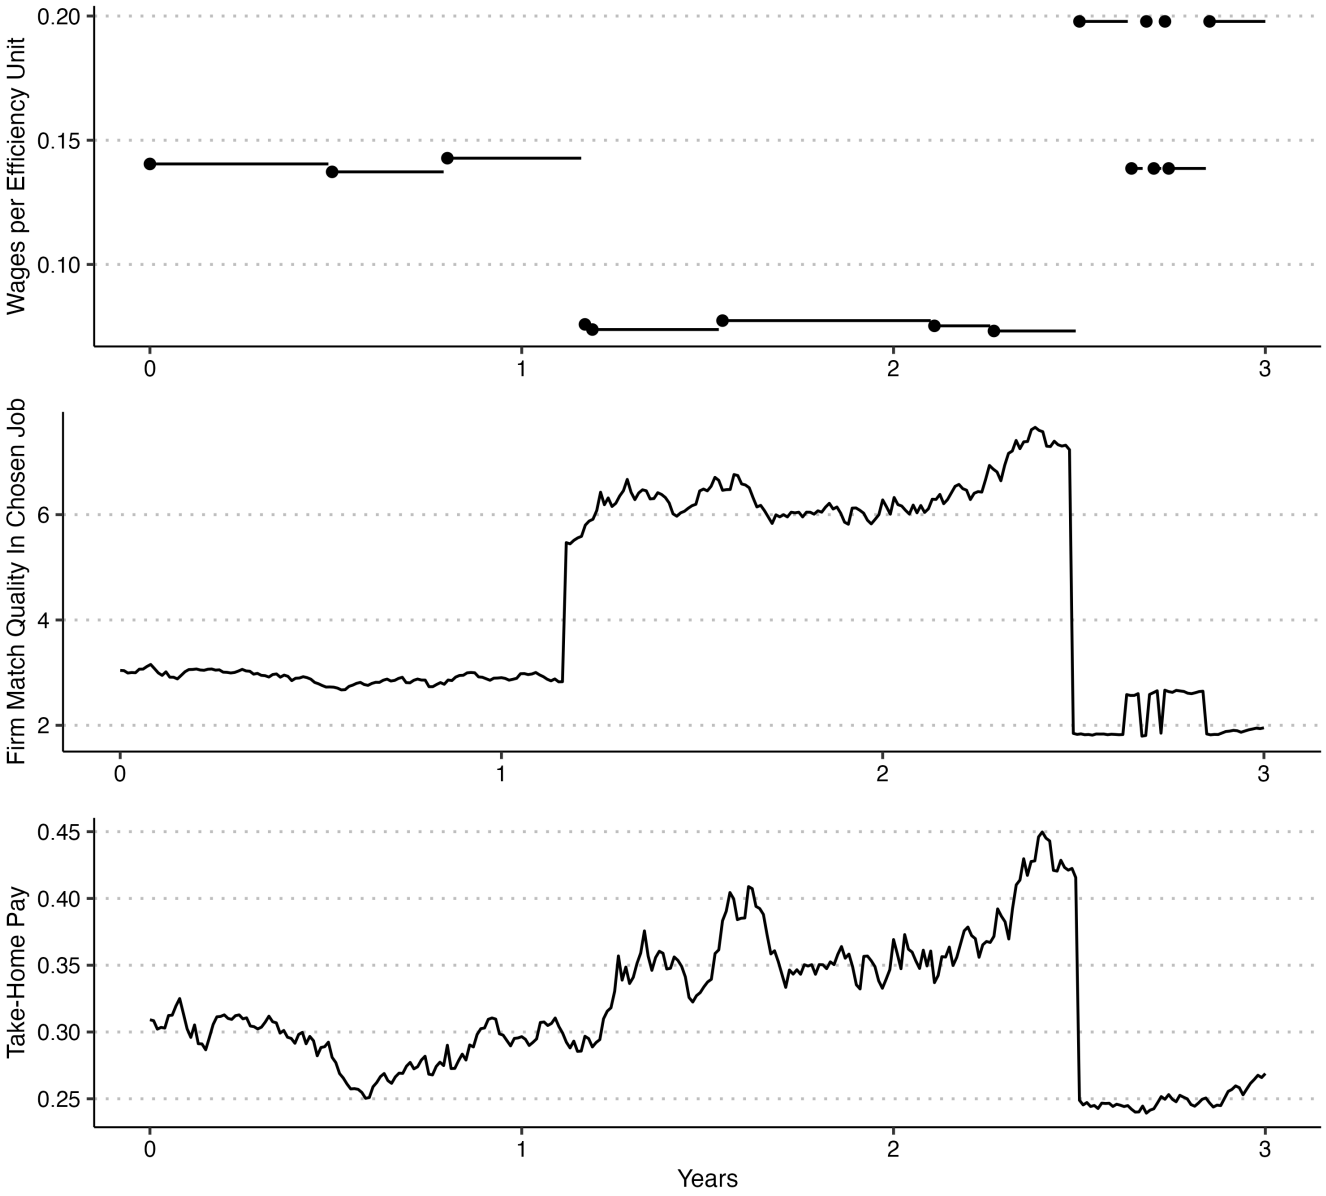
Note: Panel A of this figure shows how time for the innovation process, indexed by $\Gamma_j(t)$, relates to the time index for a person. Innovation occurs in bursts, where a shock of size μ happens instantaneously from the perspective of a worker. We set the rate of shocks to $d\Gamma_j$ to be constant over time, and occurring at rate $\frac{1}{\mu}$ so that, on average, the time indexes align. Panel B shows how a sample path for the state variable evolves over time indexed by $d\Gamma_j$. The vertical dotted lines correspond to the states that are seen by the external observer, while the continuous axis represents the unfolding of the entire process. Recall that $\hat{Q}_j(v,t) = \frac{Q_j(v,t)}{\bar{Q}_j(v,t)}$, and so the growth rate of $\log \hat{Q}_j(v,t)$ sinks as rate $\tilde{g}d\Gamma_j(t)$ conditional on $dGamma_j(v,t) = \mu$, which is the aggregate growth rate of industry quality $\bar{Q}_j(v,t)$ multiplied by the time-step. Thus, relative firm quality sinks at a constant growth rate, as competitors catch up to the firm, but the firm is able to catch up with innovation jumps of size $\log \lambda$. Panel C translates the process for innovation back to the regular time scale, showing how the two Poisson processes interact. Note that the intersections of the vertical dotted lines and the sample paths for $\log \hat{Q}$ in Panel B correspond to the observed levels of $\log \hat{Q}$ in Panel C.

Figure A.5: Static Equilibrium: Inverse Labor Supply Elasticities, Employment, and Wage Bill by Type



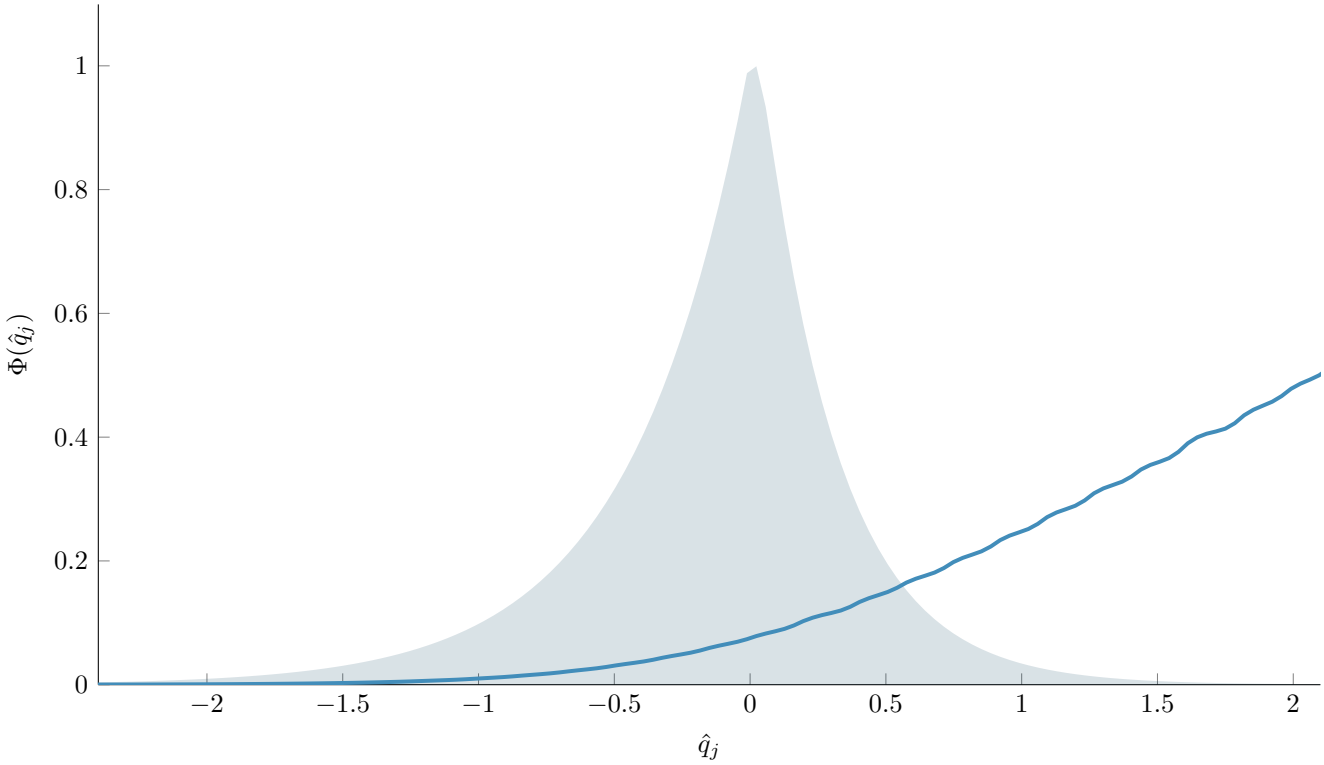
Note: Panel A of this figure shows the share of workers who are managers in terms of both wage bill and employment over the $\hat{Q}_j(v,t)$ distribution. Panel C shows the elasticities of wages with respect to the supply of labor for both types, as well as the wage bill weighted average, referred to as total in the figure. Panel C shows firm profitability over the $\hat{Q}_j(v,t)$ distribution compared to the CES-implied markup.

Figure A.6: Sample Path for a Single Worker in the Model



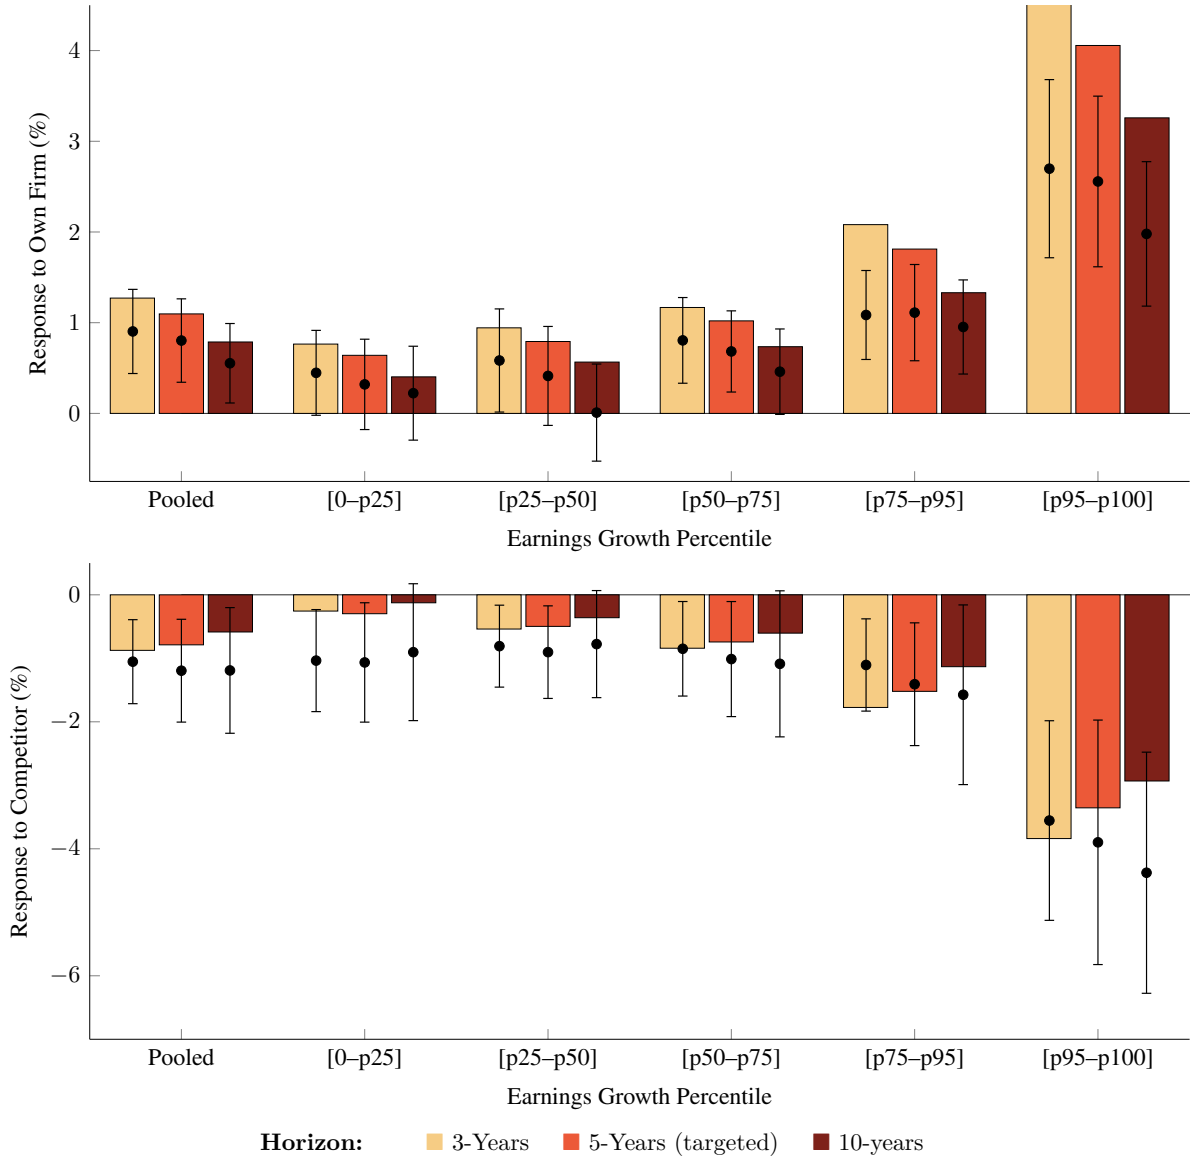
Note: This figure shows a sample path for a single manager in our model over three years. The top figure shows the wage rate associated with a worker's best job offer, which is the one that maximizes their take-home pay. The second panel shows the match quality of that job, z. The bottom panel shows how these two numbers, the efficiency wage and the match quality, aggregate to the take-home pay of the worker. Wages per efficiency unit vary through the first year because of internal improvements (and lack thereof) in the employing product line. The worker switches jobs after the first year, because they find one with a high enough match quality, z, that the take-home earnings is worth taking the cut in per-efficiency-unit pay. However, in the middle of year two that product line is displaced by a competitor, which causes them to draw a new job offer (and associated z), and re-optimize. Unfortunately for this worker, since they had such a high match quality in their employer at the start of year 2, that shock is very bad for them, which we can see in bottom panel; take-home pay drops suddenly by a little less than 50%. Subsequently, the worker is switching a few times between two jobs in her consideration set K that offer similar take-home pay.

Figure A.7: Implied Cost Function for Research and Development



Note: The solid blue line plots the implied cost function for research and development that keeps the innovation rate stationary across the $\hat{Q}$ distribution. The dashed line shows the ergodic distribution of firm (relative) productivity $\hat{q}$.

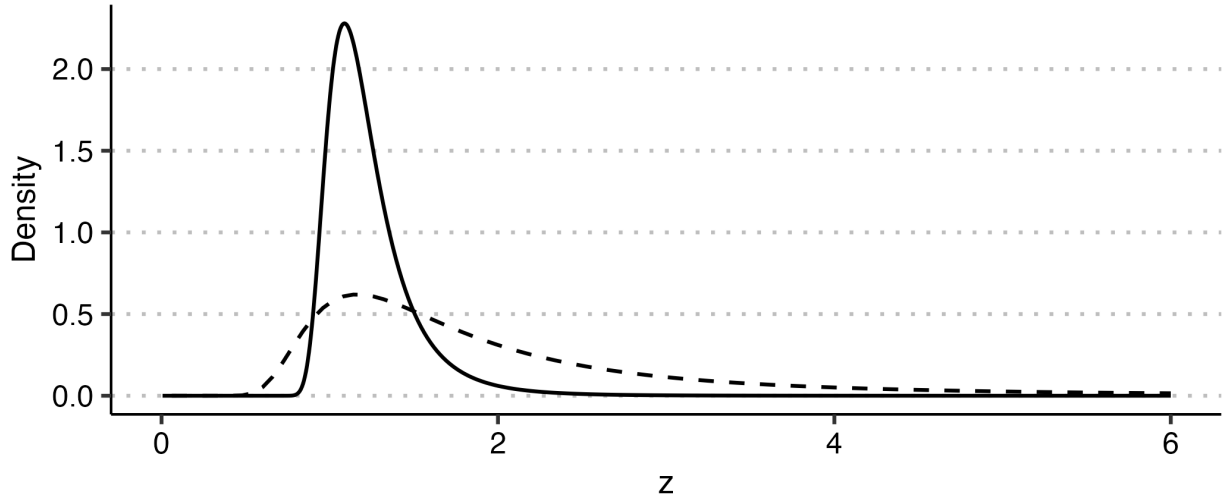
Figure A.8: Worker Wage Growth Responses Over 3, 5, and 10-Year Horizons: Model vs. Data



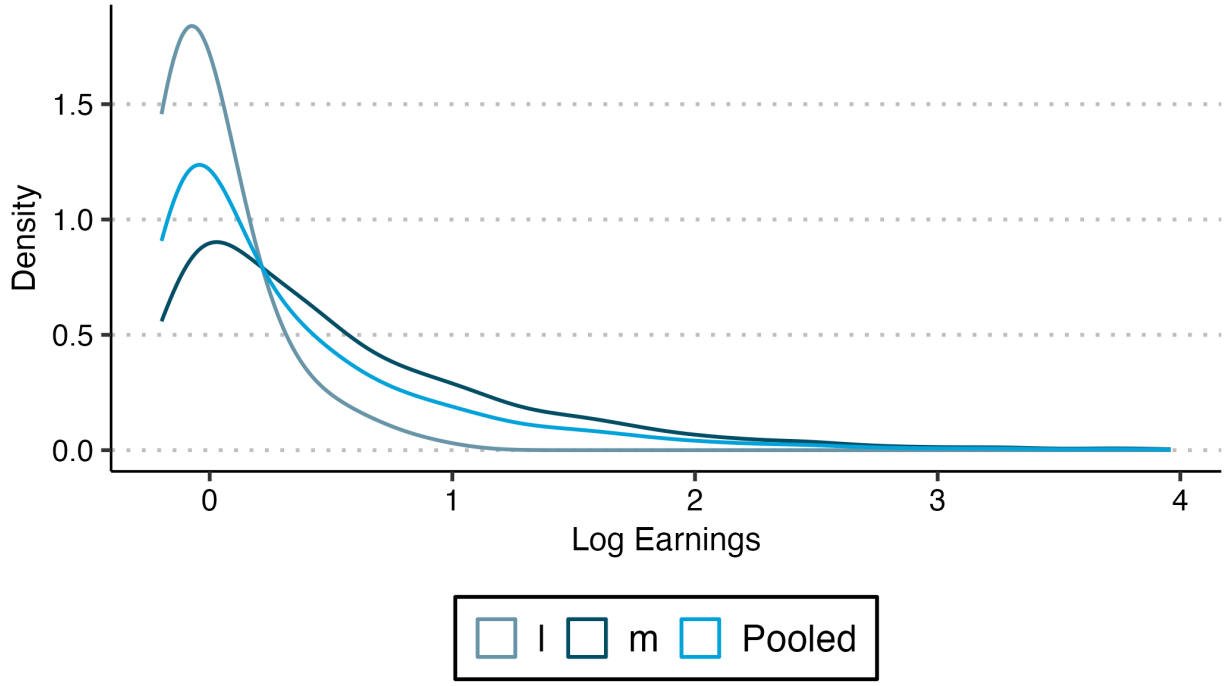
Note: This figures shows model-implied impulse responses of wage growth (columns) to product line shocks for workers both for all workers pooled together and by income percentile over 3, 5, and 10 year horizons to their empirical counterparts (points with error bars). The competitor shock includes both drops in relative firm quality $\hat{Q}$ coming from growth of competing firms, and displacement of the whole product line by a competitor. Own firm shocks come from jumps in quality within the worker's product line. Note that the 5-year responses are explicitly targeted in our calibration, while the 3 and 10 year responses are not.

Figure A.9: Distributions of Worker Skill and Earnings in the Model

Panel A: The Distribution of Human Capital (z_{ijt}) for Managers (dashed) and Laborers (whole)

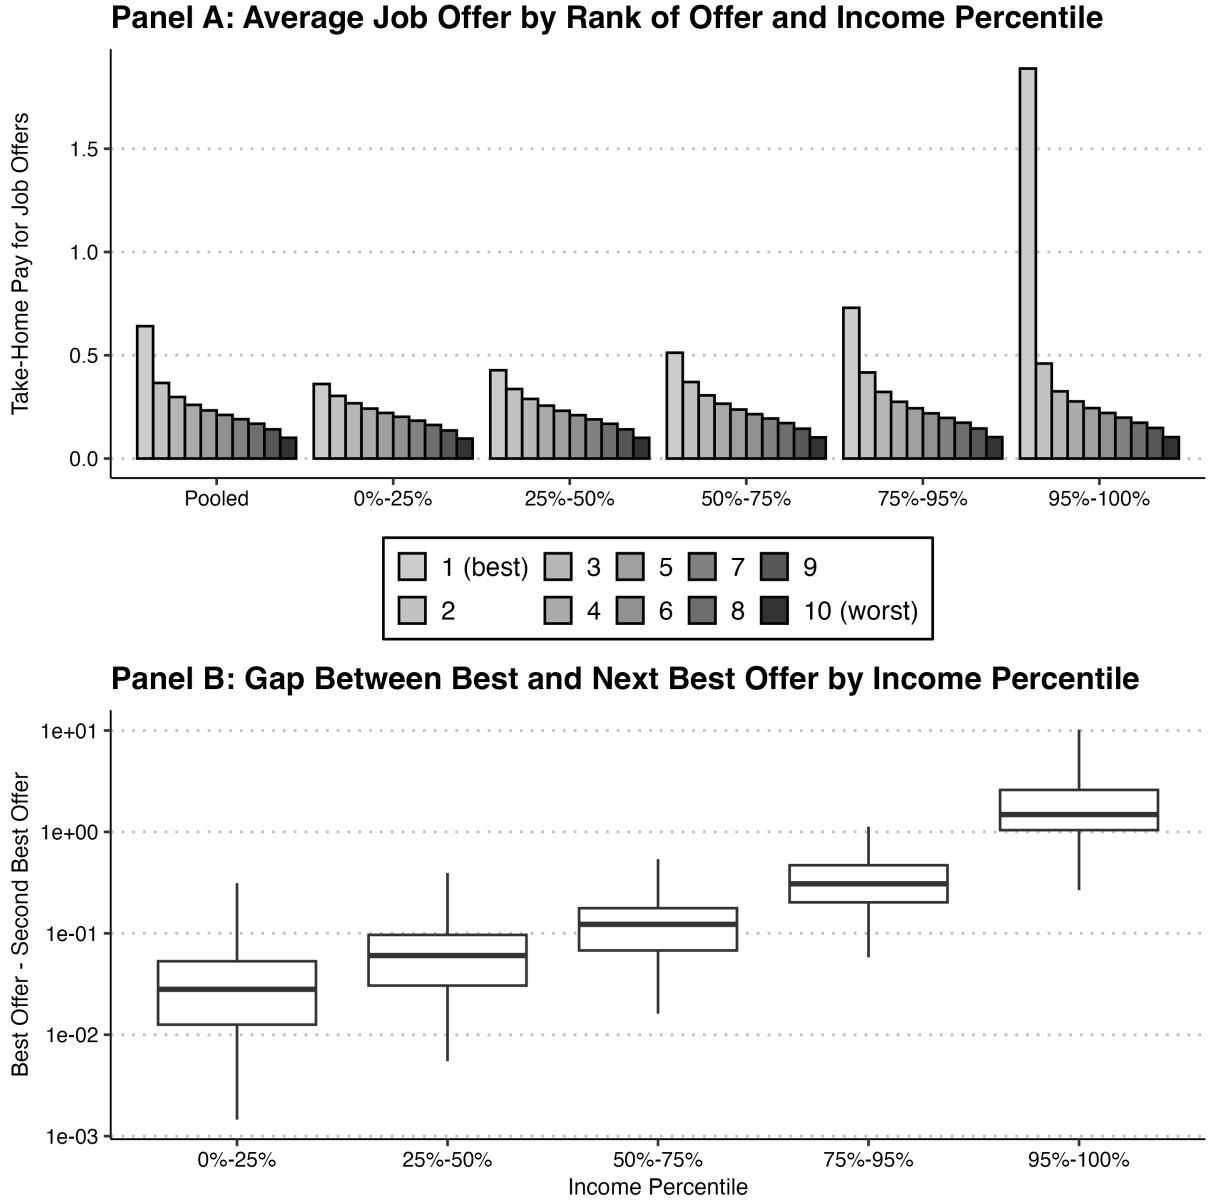


Panel B: The Distribution of Earnings for Employed Managers (m), Laborers (l), and Pooled



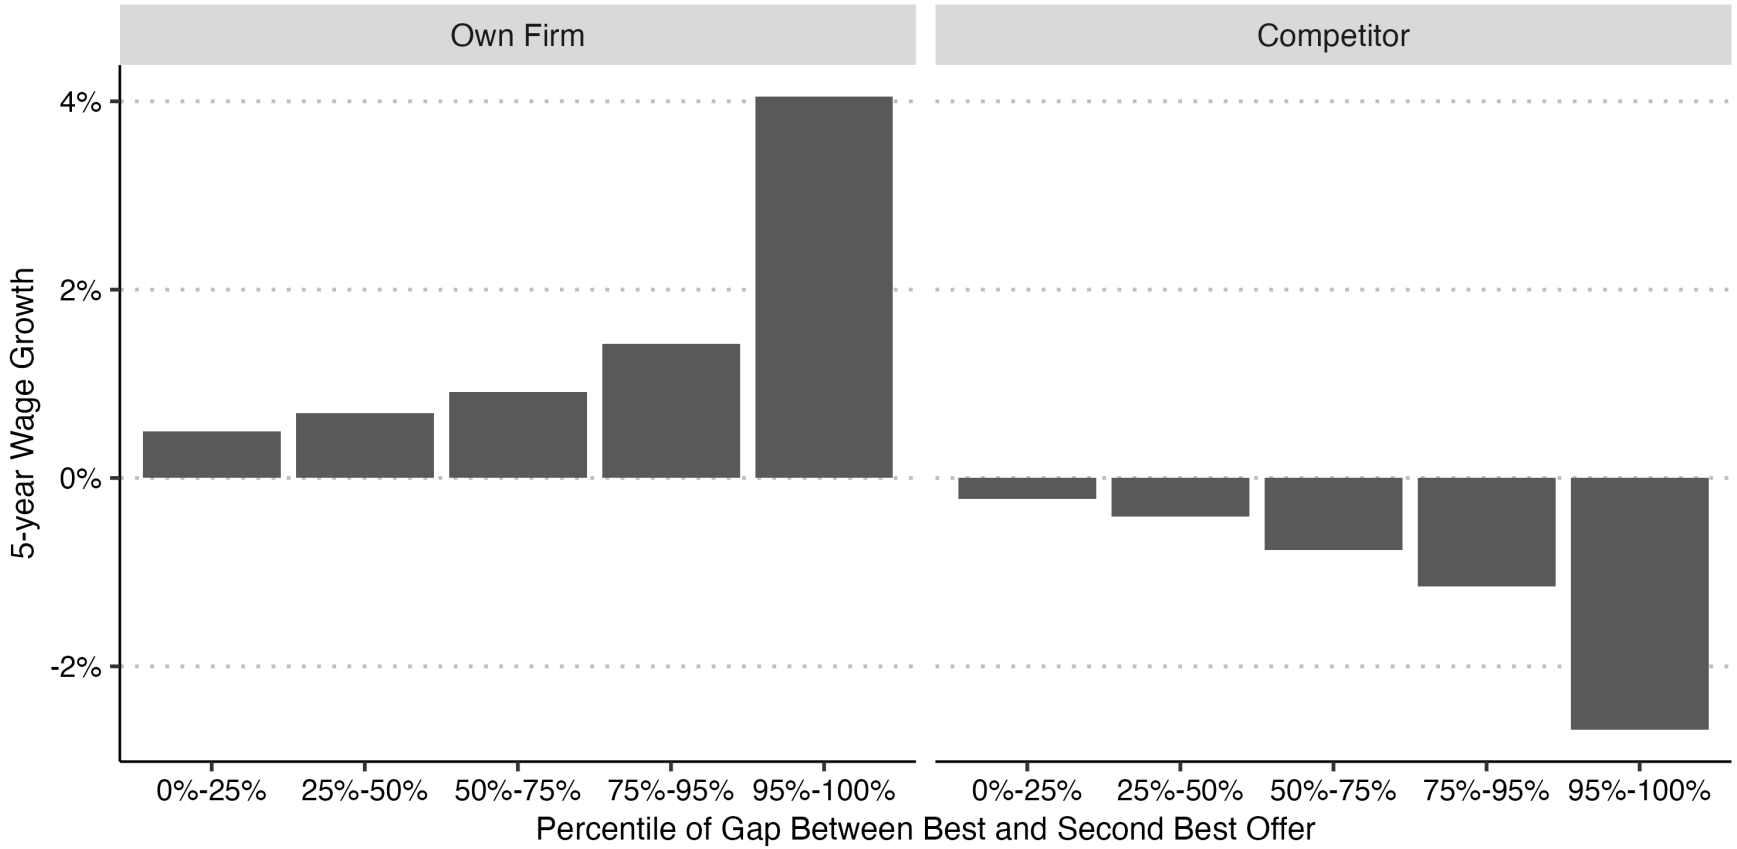
Note: Panel A of this figure shows the distribution of human capital draws, z_{ijt}, for both types. Both types have the same location parameter for the Frechet distribution, $T = 2.03$, but different dispersion parameters. For managers, $\theta_m = 2.13$ and for laborers $\theta_l = 6.81$. Panel B of this figure shows the distribution of log earnings for Managers (m), Laborers (l), and both types pooled together.

Figure A.10: Distribution of Job Offers by Worker Income Group



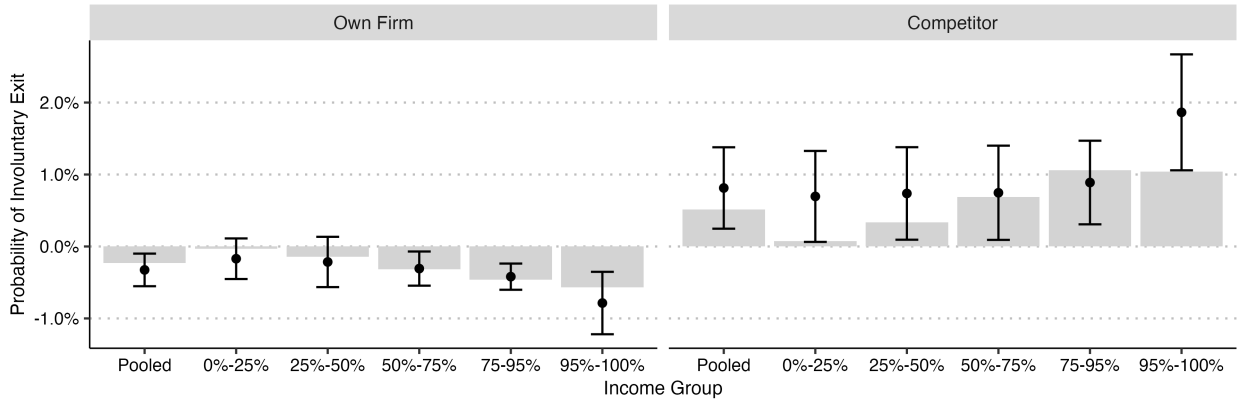
Note: This figure shows the distribution of the average job offer set by income. Jobs are ranked high (1) to low (5) and averaged within income groups. High income people have substantially worse outside options relative to their current employer, while lower income people tend to have relatively similar job offer sets in the model. Panel A shows the average job offer by income percentile and the rank of that job offer, while panel B shows boxplots of the distribution by income percentile for that gap. Note that the y axis in panel B is on a log scale.

Figure A.11: Wage Growth Responses to Innovation by Outside-Option Gap



Note: This figure shows the difference in wage growth responses to innovation shocks by percentile of the gap between a worker's best option and their second best option. Workers who have substantially better first-best options see much larger impacts; this is especially concentrated in the very top of the distribution.

Figure A.12: Probability of Involuntary Exit Following Shock



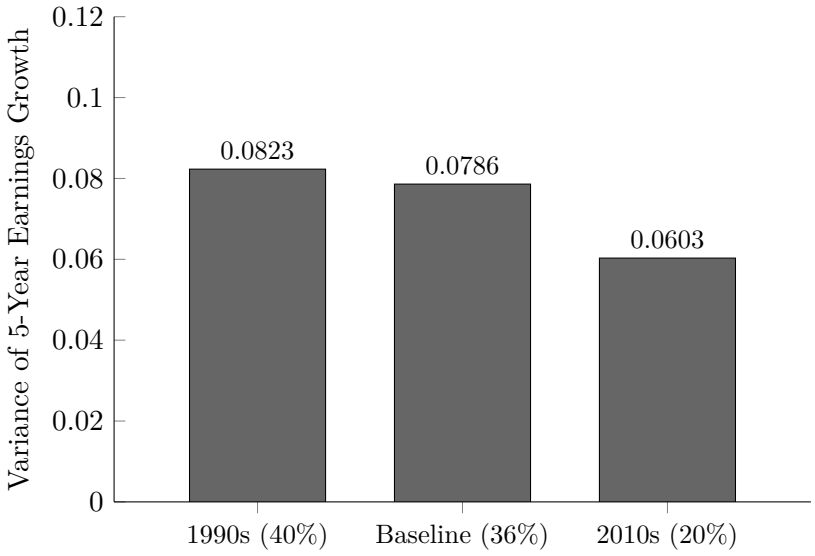
Note: Figure shows the probability of an involuntary exit, defined as having exited the firm and also having experience wage growth within the bottom 20% of the sample, in both the model and the data. The x-axis sorts the responses into income-bin interactions, and the y-axis shows the model (columns) vs. data (point + errorbar) estimates of a linear probability model. Model implied coeffients are lower in magnitude than the point estimates in the data, but are similar in sign and usually within the 95% CI of the estimates.

Figure A.13: Log Labor Share by Productivity Quartile: Model vs. Seegmiller (2024)



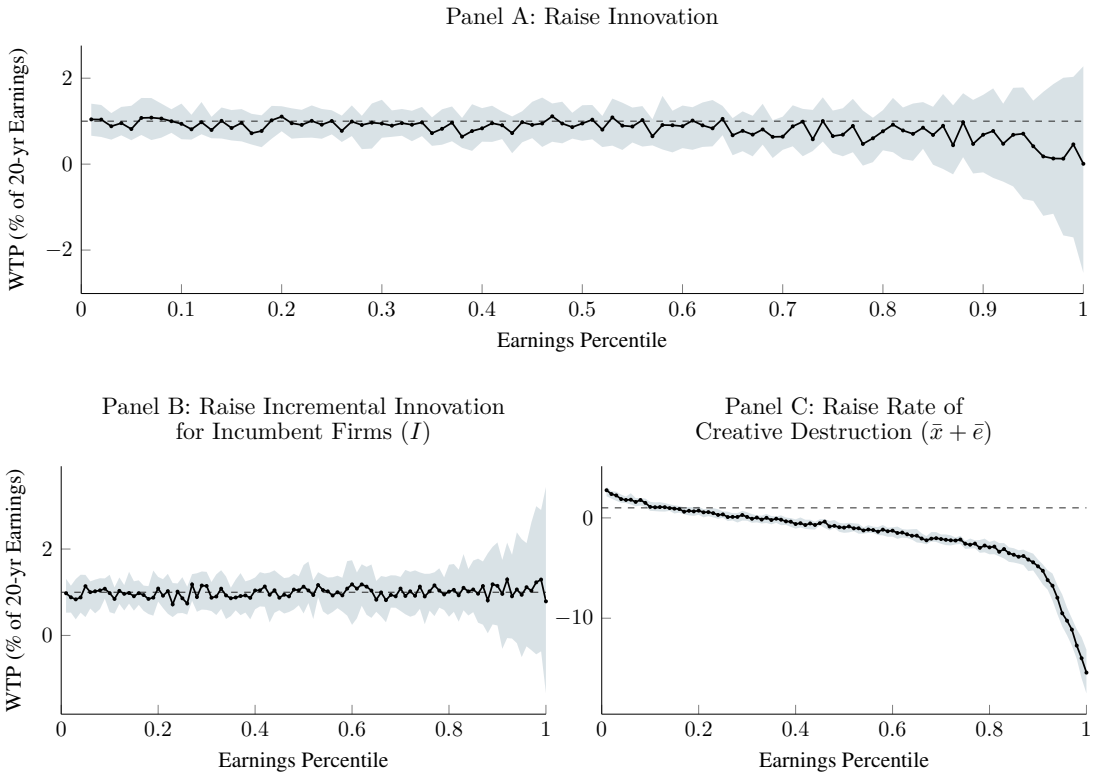
Note: Figure shows the log labor share over the same quartiles. Point estimates in the data are taken from Table 3 of Seegmiller (2024).

Figure A.14: Earnings Risk Under Declining Creative Destruction: Counterfactual Variance of 5-Year Wage Growth



Note: Variance of five-year earnings growth implied by the calibrated model under three parameterizations of the radical-innovation share. The 1990s and 2010s calibrations set the share of external innovation to match the average radical-innovation share observed in each decade (40 percent and 20 percent, respectively). The baseline calibration uses the full-sample average of 36 percent. All other parameters are held at their baseline values.

Figure A.15: Willingness to Pay for a Temporary Increase in Innovation, $\gamma = 2$



Note: Willingness-to-pay (WTP) by starting earnings percentile for temporary innovation-policy shocks under $\gamma = 2$. Panel A raises both incremental and creative-destruction innovation rates; Panel B raises incremental innovation only; Panel C raises creative destruction only. Lines show medians and shaded bands show interquartile ranges; the dashed line is the complete-markets benchmark. Specification details follow Section 4.3 and Appendix B.1.

Table A.1: Summary Statistics for Firm and Worker Samples

	Mean	5%	25%	Median	75%	95%
Panel A: Pooled Worker Sample Summary Stats						
Age	39.41	27	33	39	46	53
Wage	71,790	16,770	35,630	56,170	84,350	160,600
5-Year Log Wage Growth	-9.47%	-99.98%	-17.59%	-0.49%	13.85%	50.13%
Panel B: 75-95% Percentile Earners Sample Summary Stats						
Age	38.48	27	32	38	45	52
Wage	103,400	45,040	72,360	95,590	123,600	188,400
5-Year Log Wage Growth	-7.08%	-70.31%	-15.10%	-1.06%	10.98%	37.65%
Panel C: Top 5% of Earners Sample Summary Stats						
Age	40.87	27	35	41	47	53
Wage	263,400	80,470	127,800	177,100	268,400	636,700
5-Year Log Wage Growth	-4.44%	-81.68%	-18.17%	1.09%	19.17%	59.50%
Panel D: Firm Sample Summary Stats						
Payroll, Thousands	282,600	641	6,002	26,880	132,100	1,173,000
Employment, Thousands	8.203	0.024	0.195	0.93	4.2	35.7
Profits, Millions	532.2	1.264	10.55	45.25	214.7	1951

Note: This table contains summary statistics for both our firm and worker sample. The worker statistics are based on a linked dataset between the detailed earnings record (DER) and the current population survey (CPS). The firm data is constructed from links between Compustat and the Longitudinal Business Database (LBD).

Table A.2: Robustness Checks for Pooled OLS Results

	Own Firm Innovation		Competitor Innovation	
	Estimate	SE	Estimate	SE
A. Firm Profits				
Baseline	5.664	0.571	-4.111	1.164
Enterprise Value Denominator	5.258	0.420	-2.781	0.930
Hoberg-Phillips Competitor	5.346	0.609	-4.528	0.777
Residualized Own-Firm Innovation	4.888	0.604	-4.404	1.265
B. Firm Total Wage Bill				
Baseline	4.939	0.461	-3.428	0.921
Enterprise Value Denominator	3.755	0.371	-3.17	0.950
Hoberg-Phillips Competitor	4.766	0.483	-4.833	0.688
Residualized Own-Firm Innovation	3.647	0.479	-4.056	1.059
C. Employment				
Baseline	2.979	0.400	-2.791	0.773
Enterprise Value Denominator	1.735	0.375	-2.967	0.919
Hoberg-Phillips Competitor	2.873	0.387	-4.756	0.750
Residualized Own-Firm Innovation	2.006	0.375	-3.480	0.872
D. Pay per worker				
Baseline	1.960	0.610	-0.637	1.202
Enterprise Value Denominator	2.020	0.528	-0.203	1.322
Hoberg-Phillips Competitor	1.893	0.619	-0.077	1.018
Residualized Own-Firm Innovation	1.641	0.608	-0.576	1.372
E. Individual Worker Earnings				
Baseline	0.803	0.235	-1.195	0.413
Enterprise Value Denominator	0.521	0.184	-0.843	0.302
Hoberg-Phillips Competitor	0.858	0.261	-1.011	0.256
Residualized Own-Firm Innovation	0.639	0.179	-1.171	0.412
Adjusted Gross Income	0.866	0.210	-1.491	0.274
Wage and Salary Income	0.683	0.216	-1.337	0.352

Note: This table shows a series of robustness checks for the results in table 3. All coefficients and standard errors are multiplied by 100. "Residualized Innovation" projects out any predictable differences in the base rate of innovation using a poisson regression that has the innovation measure on the left hand size, and all of our controls on the left hand side, including industry and year fixed effects. The Hoberg-Phillips competitor definition uses the 10k textual similarity measure of Hoberg and Phillips to construct industries (using the TNIC3 definition), and computes our competitor innovation measures using this definition of industry instead of NAICS4. The Enterprise value denominator uses enterprise value instead of book assets to normalize the value of a patent filing to a firm, and baseline is the measure reported in the first figure. Adjusted Gross Income and Wage and Salary Income are different sources of wages, reported on 1040 filings, used to construct income growth rates.

Table A.3: Wage Growth Responses to Innovation by Income Group: Robustness Checks

A. Own Firm Innovation	[0% - 25%)	[25% - 50%)	[50% - 75%)	[75% - 95%)	[95% - 100%]
Baseline	0.320%	0.413%	0.683%	1.111%	2.557%
	(0.254%)	(0.278%)	(0.228%)	(0.271%)	(0.480%)
Hoberg-Phillips Competitor	0.349%	0.449%	0.787%	1.152%	2.548%
	(0.280%)	(0.291%)	(0.235%)	(0.308%)	(0.525%)
Income Rank Firm Adjusted	0.698%	0.494%	0.710%	1.239%	2.185%
	(0.304%)	(0.235%)	(0.207%)	(0.278%)	(0.483%)
Adjusted Gross Income	0.197%	0.375%	0.520%	1.011%	2.424%
	(0.215%)	(0.136%)	(0.213%)	(0.240%)	(0.497%)
Wage and Salary Income	0.314%	0.276%	0.276%	0.895%	1.823%
	(0.286%)	(0.210%)	(0.243%)	(0.259%)	(0.457%)
B. Competitor Innovation					
Baseline	-1.065%	-0.903%	-1.012%	-1.408%	-3.898%
	(0.480%)	(0.372%)	(0.462%)	(0.493%)	(0.982%)
Hoberg-Phillips Competitor	-0.435%	-1.021%	-0.848%	-1.256%	-2.99%
	(0.338%)	(0.265%)	(0.302%)	(0.414%)	(0.978%)
Income Rank Firm Adjusted	-1.771%	-0.801%	-1.047%	-0.936%	-2.214%
	(0.405%)	(0.410%)	(0.468%)	(0.485%)	(0.956%)
Adjusted Gross Income	-0.875%	-1.051%	-1.587%	-2.223%	-4.301%
	(0.220%)	(0.254%)	(0.345%)	(0.461%)	(0.897%)
Wage and Salary Income	-0.798%	-1.062%	-1.386%	-1.917%	-3.808%
	(0.345%)	(0.364%)	(0.434%)	(0.459%)	(0.839%)

Note: This table shows a series of robustness checks for the results in figure 4. "Income Rank Firm Adjusted" adjusts income by average firm pay before ranking individuals within industry when constructing the income rank percentiles for our interactions and fixed effects. The Hoberg-Phillips competitor definition uses the 10k textual similarity measure of Hoberg and Phillips to construct industries (using the TNIC3 definition), and computes our competitor innovation measures using this definition of industry instead of NAICS4.

Table A.4: Model Parameters

Description	Parameter	Value
A. Innovation Parameters		
Incremental Jump Rate	I	1.73
Displacive Innovation Rate	τ	0.084
Incumbent External Innovation Rate	x	0.063
Innovation Jump Size	$\log \lambda$	0.038
Depreciation Rate	δ	0.050
Probability a τ shock resets industry lead	p_r	0.049
B. Production Parameters		
CES Markup Parameter	σ	6.08
Weight on Managers within CES Production Function	B	0.24
Elasticity parameter for CES Production Function	ψ	0.18
Skill bias of technology	α	0.58
Share of workers who are laborers	ω	0.81
C. Worker Parameters		
Dispersion Parameter for Manager Skill Distribution	θ_m	2.13
Dispersion Parameter for Laborer Skill Distribution	θ_l	6.81
Effort Cost Parameter	φ	2.38
Value of Outside Option	R	0.433
Human Capital Persistence	ρ	0.97
D. Fixed Parameters		
Number of Job Offers	K	10
Time-Step Given Shock	μ	0.5
Cost function parameter	ζ	2

Note: Table reports the parameter estimates used in our model solution. Table 4 reports the targeted moments along with model fit.

B Additional Details of the Empirical Analysis

B.1 Data Construction

Our data on worker earnings are based on individual records tracked by the Current Population Survey (CPS) merged with employer-employee matched administrative from the Social Security Administration's Detailed Earnings Record (DER). The CPS/DER data have a panel structure, which allows us to track individuals over time and across firms and covers the 1980–2020 period. These data include information on income and employer identification numbers from Form W-2. We match the CPS/DER sample to Compustat firms using Employer Identification Numbers (EIN). The matched sample includes approximately 9,700 firms and 61.6 million worker-year observations.

We limit our sample to people ages 25 to 55 to capture people of prime working age. We also exclude self-employed workers and individuals with earnings below a minimum threshold—the amount a worker would earn working 20 hours per week for 13 weeks at the federal minimum wage following Guvenen et al. (2014). After restricting the sample to observations for which we have sufficient income history and firm controls (for instance, lagged levels of R&D spending), we are left with 12,500 firms and 132.000 firm years in the firm panel, and 9,700 firms, 780,000 people, and 5.6 million person years in the worker panel.

Table A.1 shows summary statistics for our firm and worker samples. Panel A contains information on the pooled sample, while B and C condition on the income groups we focus on when we calibrate our model. Wage growth has large tails: the 5% and 95% quantiles are substantially larger than even the 25% and 75%. Wealthier workers do tend to be slightly older, which we can see comparing panels B and C, which makes including age and lifecycle controls in our regressions important. Panel D contains information on our firm sample. Firms in our sample are all public, due to the requirement that we are able to observe market reactions to patent announcements, but have substantial heterogeneity in size.

The firm wage bill and employment comes from the Longitudinal Business Database (LBD), a longitudinal dataset of non-farm firms at the establishment level. Jarmin and Miranda (2002) and Chow, Fort, Goetz, Goldschlag, Lawrence, Perlman, Stinson, and White (2021) describe the construction of the dataset in substantial detail, and we refer readers to these documents. For our purposes, the LBD provides extremely detailed wage and employment which is sourced from tax filings. We link these records with Compustat, which we use for firm control variables as well as firm profits.

To construct a proxy for the level of pay to current employees, we take the individual records from the DER-CPS linked sample we use for the worker-level regressions. The focal firm is taken as the firm employing them at $t = 0$. To track the changes in wages paid to employees at the time of the innovation, we look at the wages paid to that set of employees by the focal firm at $t \in \{-5, ..., 10\}$,

and track that forward and backwards. We then aggregate those numbers by firm, and calculate the growth of that number in the same fashion as our other firm level variables, running firm-level regressions with the same controls as the regressions in Figure 1. This variable therefore omits the compensation of these incumbent workers who subsequently leave the firm, or the compensation of subsequently hired workers.

Worker level regressions We construct flexible controls for worker age and lagged earnings by linearly interpolating between 3rd degree Chebyshev polynomials in workers' lagged income quantiles within an industry-age bin at 10-year age intervals. In addition, to soak up some potential variation related to potential mean-reversion in earnings (which could be the case following large transitory shocks), we also include 3rd degree Chebyshev polynomials in workers' lagged income growth rate percentiles, and we allow these coefficients to differ across five bins formed based upon a worker's income rank within their firm's industry. To ensure that our point estimates are comparable to the analysis in Section 2.2, in which the unit of observation is at the firm-year as opposed to the worker-year level, we weigh observations by the inverse of the number of workers in each firm-year. Standard errors are clustered at the industry-year level. For our quantile regression specifications, we compute standard errors using a block-resampling procedure that allows for persistence at the sic3 industry level (the analogue of clustering by industry).

A primary point of focus in the paper is the concentration of the innovation risks (via profit sharing) on highly paid employees. We construct a measure of income percentile based on a worker's wage information reported in the Detailed Earnings Record relative to their NAICS4 industry within the year we observe. Income bins are indicators based on a workers current income percentile within their industry, broken into [0%, 25%), [25%, 50%), [75%, 95%), and [95%, 100%] bins.

The vector $\mathbf{Z}$ includes several controls, including one lagged value of the dependent variable and the log of the book value of firm assets and the firm's market capitalization to alleviate our concern that firm size may introduce some mechanical correlation between the dependent variable and our innovation measure.

We also include controls for the depreciated stock of prior patents and spending on research and development as a portion of firm assets, since firms with prior valuable innovations and high R&D spending may persistently innovate at a different rate. In order to construct a patent stock, we depreciate past innovations at a rate of 20% per year, with the current patent stock being equal to the historical sum of a firm's patents, weighted with a discount rate of 0.8. We also control for firm idiosyncratic volatility σ_{ft} because it may have a mechanical effect on our innovation measure and is likely correlated with firms' future growth opportunities or the risk in worker earnings. Additionally, we control for firm profitability, defined as sales minus cost of goods sold, scaled by assets.

We construct controls for worker age and lagged earnings by linearly interpolating between 3rd degree Chebyshev polynomials in workers' lagged income quantiles within an industry-age

bin at 10-year age intervals. In addition, to soak up some potential variation related to potential mean-reversion in earnings (which could be the case following large transitory shocks), we also include 3rd degree Chebyshev polynomials in workers' lagged income growth rate percentiles, and we allow these coefficients to differ across five bins formed based upon a worker's income rank within their firm's industry.

Firm and Industry Splits For the exercises in Table 1, we split firms into high / low groups relative to median along a number of dimensions. The industry innovation rate is measured by the value of the patents an industry puts out each year. The firm patent stock relative to industry median is based on the accumulated depreciated stock of patents described above. Income inequality is based on the annual cross-sectional standard deviation of average firm wages in a given year. Average industry wages are measured using our worker-level CPS / DER linked sample.

To compute firm κ data analogues that we can compare to the model, we use the panel of firms in Compustat and for each year compute the relative average wage, profits, and depreciated patent stock of the firm relative to its NAICS4 industry. We then run principal components analysis on the panel, and extract the first principal component as our data proxy for κ.

For industry variables we compute the median level of the measure by industry, grouping industries into high if they are above the median and low otherwise. In the case of the firm variables, we perform this sort within industry relative to the median value of a firm's patent stock within a NAICS4 industry.

Robustness Measures Our robustness checks include an alternative text-based definition of industry rather than NAICS4. We borrow the Hoberg and Phillips (2016) measure of product market competition, which is based on the language similarity of filings in form 10k, and perform the same analyses as 12 and 15. The other main robustness check we takes account of potential mis-measurement of firm size. Rather than using the market capitalization of the firm as the normalizing factor, K, in equation 10. We define enterprise value as market capitalization plus book assets minus the value of common equity. Finally, we construct a "residualized" measure of innovation, which accounts for potential firm-level predictability in the KPSS measure. To construct this, we estimate a Poisson regression that conditions on all of the firm-level controls in $\mathbf{Z}$ and industry and year fixed effects. The difference between the realized KPSS value and the value predicted by this regression is the residualized measure we substitute for A_f in our regressions in table A.3. Because there may be income sources that do not appear on form W-2, we repeat our analysis with two forms of income reported on form 1040: Adjusted Gross Income (AGI) and Wage and Salary Income (WSI). Data is available only intermittently from the period 1990-2000, when it becomes a continuous panel. For years that are missing information, the three year backward and five year forward average earnings are computed using the average of reported earnings available

for the relevant periods. The same participation restrictions present in the DER worker panel are applied to the 1040 panel.

C Model Appendix

In this section, we describe the setup and equilibrium conditions of our model in greater detail, as well as other implementation details. For completeness, we reproduce some key equations from the main text here. Table 5 summarizes the key notation used throughout the model.

Table 5: Key Model Notation

Symbol	Description
Aggregate and Industry	
$Y(t)$	Aggregate output
$Y_j(v,t)$	Output of variety v in industry j
$\bar{Q}_j(t)$	Average productivity in industry j
$\hat{Q}_j(v,t)$	Relative productivity of variety v: $Q_j(v,t)/\bar{Q}_j(t)$
σ	Elasticity of substitution across varieties
$\Gamma_j(t)$	Industry-level business time (subordinator)
μ	Innovation burst size parameter
Innovation	
I	Internal innovation intensity
x	External (business-stealing) innovation intensity
e	Entry rate of new firms
λ	Step size of productivity improvement
δ	Depreciation rate of productivity
$\tilde{g}$	Growth rate of industry productivity
Production and Labor	
$l_j(v,t)$, $m_j(v,t)$	Efficiency units of laborers and managers
$w_j(v,t)$, $s_j(v,t)$	Piece rates for laborers and managers
B	Relative importance of managers in production
ψ	Elasticity of substitution between labor types
α	Skill-bias parameter for managerial productivity
Workers	
$z_j(i,v,t)$	Match-specific productivity of worker i
$z^*(i,k,t)$	Transformed (normal) match quality index
θ_l, θ_m	Fréchet shape parameters for laborers and managers
φ	Disutility of effort parameter
K	Number of job options per worker
R	Outside option (home production)
ρ	Mean-reversion rate of match quality
Firm Dynamics	
$V(\hat{q})$	Value of a product line with relative productivity $\hat{q}$
$\pi(\hat{q})$	Flow profits of a product line
$N_{f,t}$	Number of product lines operated by firm f

C.1 Static Equilibrium Conditions

Relative Prices

There is a single consumption good $Y(t)$ produced by a perfectly competitive final goods sector that is a nested CES aggregate of intermediate goods, $Y_j(v,t)$, indexed by variety v and industry j:

$$Y(t) = \exp\left(\int_0^1 \log\left(\int_0^1 Y_j(v,t)^{\frac{\sigma-1}{\sigma}} dv\right)^{\frac{\sigma}{\sigma-1}} dj\right), \tag{C.1}$$

where $\sigma > 1$ is the elasticity of substitution between product varieties within an industry. The production function for intermediate good in industry j and variety v is produced only by the most efficient producer according to

$$Y_j(v,t) = Q_j(v,t)\left((1-B)\left[Q_j(v,t)\, l_j(v,t)\right]^{\frac{\psi-1}{\psi}} + B\left[\bar{Q}_j(t)^{\alpha} Q_j(v,t)^{-\alpha} m_j(v,t)\right]^{\frac{\psi-1}{\psi}}\right)^{\frac{\psi}{\psi-1}}. \tag{C.2}$$

where $l_j(v,t)$ and $m_j(v,t)$ captures the number of efficiency units of laborers and managers, respectively, and ψ is the elasticity of substitution between skill types. We let

$$\bar{Q}_j(t) \equiv \int_0^1 Q_j(v,t)dv \quad \text{and} \quad \hat{Q}_j(v,t) \equiv \frac{Q_j(v,t)}{\bar{Q}_j(t)} \tag{C.3}$$

denote the average productivity across all varieties within industry j and the relative productivity in variety v, respectively.

Using the final goods producer's input demand from the outer nest of the production function (C.1), we get that the demand for industry j output $Y_j(t) = \left(\int_0^1 Y_j(v,t)^{\frac{\sigma-1}{\sigma}} dv\right)^{\frac{\sigma}{\sigma-1}}$, where $Y_j(v,t)$, satisfies

$$Y_j(t) = \frac{Y(t)}{P_j(t)}, \tag{C.4}$$

an expression which holds for all industries $j \in [0,1]$. Within industries, the demand for intermediate good $Y_j(v,t)$ is

$$Y_j(v,t) = Y_j(t)\left[\frac{p_j(v,t)}{P_j(t)}\right]^{-\sigma} = \frac{Y(t)}{P_j(t)}\left(\frac{P_j(v,t)}{P_j(t)}\right)^{-\sigma}, \tag{C.5}$$

an expression which holds for all varieties, v, and industries $j \in [0,1]$, where the price index for industry j output, $P_j(t)$, is given by:

$$P_j(t) = \left(\int_0^1 P_j(v,t)^{1-\sigma} dv\right)^{\frac{1}{1-\sigma}}. \tag{C.6}$$

Further notice that revenue only depends on relative prices:

$$p_j(v,t)Y_j(v,t) = Y(t)\left(\frac{P_j(v,t)}{P_j(t)}\right)^{1-\sigma} \equiv Y(t)\left[\hat{P}_j(v,t)\right]^{1-\sigma},$$

a property which will help immensely along the balanced growth path.

Rearranging (C.4) and equating it for two different industries yields the following condition

$$P_j(t)Y_j(t) \equiv P_j(t)\bar{Q}_j(t)\left(\int_0^1 \left[\frac{Y_j(v,t)}{\bar{Q}_j(t)}\right]^{\frac{\sigma-1}{\sigma}} dv\right)^{\frac{\sigma-1}{\sigma}} = P_k(t)\bar{Q}_k(t)\left(\int_0^1 \left[\frac{Y_k(v,t)}{\bar{Q}_k(t)}\right]^{\frac{\sigma-1}{\sigma}} dv\right)^{\frac{\sigma-1}{\sigma}}. \tag{C.7}$$

Along the balanced growth path we consider, the distribution of $\hat{p}_j(v,t)$ and $\hat{Q}_j(v,t)$ across varieties will be constant over time, hence the integral over varieties is the same on both sides of (C.7), which simplifies to

$$\frac{P_j(t)}{P_k(t)} = \frac{\bar{Q}_k(t)}{\bar{Q}_j(t)} \tag{C.8}$$

and aggregate output satisfies

$$Y(t) = \int_0^1 P_j(t)Y_j(t)dj = \int_0^1 \bar{Q}_j(t) \underbrace{\left(\int_0^1 \left[\frac{Y_j(v,t)}{\bar{Q}_j(t)}\right]^{\frac{\sigma-1}{\sigma}} dv\right)^{\frac{\sigma-1}{\sigma}}}_{\equiv \bar{A}} dj = \bar{A}\bar{Q}(t). \tag{C.9}$$

Labor Markets and Equilibrium Wages

Workers of type $\tau \in \{l, m\}$ observe K job offers and for each job $k \in K$ have firm-specific match quality

$$z_{ik} \sim \text{Fréchet}(T_\tau, \theta_\tau) \tag{C.10}$$

Each of those jobs pays wages per efficiency unit $\{w_k\}_{k=1}^K$. In general we denote wages for type m as s, but because things are symmetric for now we will only use w in order to conserve space. Workers face an effort cost

$$c(e) = \tfrac{e^{1+\phi}}{1+\phi},\ \phi > 0 \tag{C.11}$$

which implies optimal effort

$$e^\star_{ik} = (w_k z_{ik})^{1/\phi},\ \beta = 1 + \tfrac{1}{\phi}. \tag{C.12}$$

This gives total earnings per unit of output

$$y_{ik} = w_k e^\star_{ik} = (w_k z_{ik})^\beta = w_k^\beta z_{ik}^\beta. \tag{C.13}$$

A power function of a Fréchet distributed random variable is still Fréchet distributed, so we know

that the distribution of take-home earnings offered in closed form. Further, the maximum of a series of Fréchet random variables is also Fréchet, with distribution

$$z_{ik}^{\beta} \sim \text{Fréchet}(T_\tau, \tilde{\theta}_\tau), \quad y_{ik} \sim \text{Fréchet}(S_k, \tilde{\theta}_\tau), \qquad S_k = T_\tau w_k^{\theta_\tau} \;\; (\beta\tilde{\theta}_\tau = \theta_\tau). \tag{C.14}$$

So adding the effort choice only impacts the shape parameter but leaves the scale parameter unchanged.

Let $\Lambda_\tau = T_\tau \sum_{k=1}^{K} w_k^{\theta_\tau}$. Then the distribution of the maximum of K offers is also distributed Fréchet,

$$y_i^* = \max_{k \le K} y_{ik} \sim \text{Fréchet}(\Lambda_\tau, \tilde{\theta}_\tau) \tag{C.15}$$

However, all workers also face a common outside option which pays wages at rate R. All workers have this option, and have a single efficiency unit in that option. This means that their post-effort wage rate is R^β. Their probability of participating in the market is given by

$$\pi_\tau = 1 - \exp(-\Lambda_\tau R^{-\theta_\tau})$$

which follows from the CDF of the Fréchet distribution.

This gives us a final probability that the worker chooses a given job,

$$P\{\text{choose k}\} = \pi_\tau \, \frac{w_k^{\theta_\tau}}{\sum_{\kappa=1}^{K} w_\kappa^{\theta_\tau}}$$

Three-channel decomposition of labor supply. Given the extensive- and intensive-margin decisions above, the aggregate labor supply of laborers and managers faced by a firm are deterministic functions of its posted wages w and s:

$$l(w,t) = \omega \, K \, T_l^{[1+1/\varphi]/\theta_l} \, w^{\frac{1}{\varphi}} \times \mathbb{E}_{\mathcal{W}_{i,t}} \left[\left(\frac{w^{\theta_l}}{w^{\theta_l} + \mathcal{W}_i^{\theta_l}} \right)^{[1+1/\varphi]/\theta_l - 1} \gamma\left(1 - \tfrac{1+1/\varphi}{\theta_l}, -\left[T_l(w^{\theta_l} + \mathcal{W}_{i,t}^{\theta_l})\right] R^{-\theta_l}\right) \right] \tag{C.16}$$

$$m(s,t) = (1-\omega) \, K \, T_m^{[1+1/\varphi]/\theta_m} \, s^{\frac{1}{\varphi}} \times \mathbb{E}_{\mathcal{S}_{i,t}} \left[\left(\frac{s^{\theta_m}}{s^{\theta_m} + \mathcal{S}_{i,t}^{\theta_m}} \right)^{[1+1/\varphi]/\theta_m - 1} \gamma\left(1 - \tfrac{1+1/\varphi}{\theta_m}, -[T_m(s^{\theta_m} + \mathcal{S}_{i,t}^{\theta_m})] R^{-\theta_m}\right) \right] \tag{C.17}$$

where $\gamma(\cdot,\cdot)$ denotes the lower incomplete gamma function, which arises because the distribution of the maximum of truncated Fréchet draws appears in the derivation; see Nadarajah (2009) for details. The terms $\mathcal{W}_{i,t}$ and $\mathcal{S}_{i,t}$ summarize the posted wages offered by the focal firm's $K-1$ randomly

sampled competitors:

$$\mathcal{W}_{i,t} \equiv \left\{ \sum_{\kappa=1}^{K-1} [w_{j(\kappa,i,t)}(v(\kappa,i,t),t)]^{\theta_l} \right\}^{\frac{1}{\theta_l}} \quad \text{and} \quad \mathcal{S}_{i,t} \equiv \left\{ \sum_{\kappa=1}^{K-1} [s_{j(\kappa,i,t)}(v(\kappa,i,t),t)]^{\theta_m} \right\}^{\frac{1}{\theta_m}}. \tag{C.18}$$

Along a balanced growth path these functions are time-invariant. Three channels link a firm's posted wage to the quantity of labor it attracts: (i) the *intensive margin* ($w^{1/\varphi}$), reflecting workers' effort choices; (ii) *reallocation across employers* (the term inside the expectation), capturing how a higher wage draws workers away from competing firms; and (iii) *non-employment* (the incomplete gamma function), governing the inflow from home production. As a firm grows, the latter two channels saturate and the labor-supply elasticity asymptotes to $1/\varphi$, the source of endogenous monopsony power for large, productive firms.

Firm Profit Maximization Problem

Next, we consider the profit maximization problem faced by the leading producer in a given product line. We assume that the equilibrium in product markets is as in Acemoglu et al. (2012, 2018). Specifically, we assume that there is a two stage pricing game in which firms first must pay a small fixed cost $\varepsilon > 0$ in order to enter, where we take the limit as $\varepsilon \downarrow 0$. Then, in a second stage, the firms which have entered compete a la Bertrand by simultaneously posting prices. Given our CES demand assumption, in equilibrium, the most efficient producer will post a price equal to the standard CES markup over marginal cost.

The firm profit maximization problem, expressed relative to aggregate output, is:

$$\pi_j(v,t) = \max_{\hat{P}} \left[\hat{P}^{1-\sigma} - \left\{ w\left(\frac{\bar{A}}{\hat{Q}_j(v,t)} \hat{P}^{-\sigma} \right) + \hat{Q}_j(v,t)^{\alpha} s\left(\frac{\bar{A}}{B\hat{Q}_j(v,t)^{1-\alpha}} \hat{P}^{-\sigma} \right) \right\} \frac{\bar{A}}{\hat{Q}_j(v,t)} \hat{P}^{-\sigma} \right] \tag{C.19}$$

The firm trades off its product-market power (setting a high price) against the cost of attracting labor (paying higher wages). The functions $w(\cdot)$ and $s(\cdot)$ are the inverses of the labor supply curves in equation (C.16).

Properties of the static equilibrium. Appendix Figure A.5 summarizes the key features of the calibrated static equilibrium. Panel A shows how $\alpha > 0$ implies that more productive firms employ a disproportionately larger share of managers: laborers and managers are complements ($\psi < 1$) but managerial productivity does not scale one-for-one with firm productivity, so firms require a higher manager-to-laborer ratio as they expand. Panel B reports the labor supply elasticities for laborers and managers; in the calibration $\theta_l > \theta_m$, making managers' extensive-margin labor supply less elastic. For a given worker type, more productive firms face a less elastic labor supply curve because they hire more workers in total—at higher levels of productivity, firms must raise wages substantially to attract workers from other firms. Panel C shows that firm profitability increases with the productivity advantage $\hat{q}$: more productive firms face steeper labor supply curves,

must raise wages sharply to expand employment, and therefore optimally post wages further below the marginal product of labor—generating higher profit margins consistent with the evidence in Gouin-Bonenfant (2022); Seegmiller (2024). Panel D shows that average wages paid by the firm increase in firm productivity, for both laborers and managers, because firms that need more labor must offer higher wages.

$$\begin{aligned} \Pi_j(v,t) &= \max_{P,l,m} Y(t)\left[\left(\frac{P}{P_j(t)}\right)^{1-\sigma} - w(l)l - s(m)m\right], \\ \text{s.t. } Q_j(v,t)\left[(1-B)l^{\frac{\psi-1}{\psi}} + B[\hat{Q}_j(v,t)^{-\alpha}m]^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}} &= Y_j(t)\left(\frac{P}{P_j(t)}\right)^{-\sigma} = \frac{Y(t)}{P_j(t)}\left(\frac{P}{P_j(t)}\right)^{-\sigma} \end{aligned} \tag{C.20}$$

Incorporating equations (C.7) and (C.9), yields the following identity

$$\frac{Y(t)}{P_j(t)Q_j(v,t)} = \frac{\bar{A}\bar{Q}(t)}{P_j(t)Q_j(v,t)} = \frac{\bar{A}\bar{Q}(t)}{Q_j(v,t)}\frac{\bar{Q}_j(t)}{\bar{Q}(t)} = \frac{\bar{A}}{\hat{Q}_j(v,t)}, \tag{C.21}$$

where $\hat{Q}_j(v,t) \equiv \frac{Q_j(v,t)}{\bar{Q}_j(t)}$. Hence, we can write the constraint as

$$\left[(1-B)l^{\frac{\psi-1}{\psi}} + B[\hat{Q}_j(v,t)^{-\alpha}m]^{\frac{\psi-1}{\psi}}\right]^{\frac{\psi}{\psi-1}} = \frac{Y(t)}{Q_j(v,t)P_j(t)}\left(\frac{P}{P_j(t)}\right)^{-\sigma} = \frac{\bar{A}}{\hat{Q}_j(v,t)}\left(\frac{P}{P_j(t)}\right)^{-\sigma} = \frac{\bar{A}\hat{Y}_j(v,t)}{\hat{Q}_j(v,t)}, \tag{C.22}$$

where the first equality uses C.21 and the second equality uses C.5.

Next, defining $\hat{Y}_j(v,t) \equiv \frac{Y_j(v,t)}{\bar{Y}_j(t)}$, we derive the labor requirement functions:

$$\chi(\hat{Q}_j(v,t),\hat{Y}_j(v,t)) \equiv \frac{m(\hat{Q}_j(v,t),\hat{Y}_j(v,t))}{l(\hat{Q}_j(v,t),\hat{Y}_j(v,t))} = \frac{B^{\psi}\hat{Q}_j(v,t)^{\alpha(1-\psi)}}{(1-B)^{\psi}}\frac{mc_l(\hat{Q}_j(v,t),\hat{Y}_j(v,t))^{\psi}}{mc_m(\hat{Q}_j(v,t),\hat{Y}_j(v,t))^{\psi}} \tag{C.23}$$

$$l(\hat{Q}_j(v,t),\hat{Y}_j(v,t)) = \frac{\bar{A}\hat{Y}_j(v,t)}{\hat{Q}_j(v,t)}\left[(1-B) + B[\hat{Q}_j(v,t)^{-\alpha}\chi(\hat{Q}_j(v,t),\hat{Y}_j(v,t))]^{\frac{\psi-1}{\psi}}\right]^{\frac{-\psi}{\psi-1}} \tag{C.24}$$

$$mc_l(\hat{Q}_j(v,t),\hat{Y}_j(v,t)) = w(l(\hat{Q}_j(v,t),\hat{Y}_j(v,t)))(1+\tfrac{d\log w(l)}{d\log l}|_{l=l(\hat{Q}_j(v,t),\hat{Y}_j(v,t))}), \tag{C.25}$$

$$mc_m(\hat{Q}_j(v,t),\hat{Y}_j(v,t)) = s(m(\hat{Q}_j(v,t),\hat{Y}_j(v,t)))(1+\tfrac{d\log s(m)}{d\log m}|_{m=m(\hat{Q}_j(v,t),\hat{Y}_j(v,t))}), \tag{C.26}$$

$$\chi_j(v,t) \equiv \frac{m_j(v,t)}{l_j(v,t)} = \frac{B^{\psi}\hat{Q}_j(v,t)^{\alpha(1-\psi)}}{(1-B)^{\psi}}\frac{[w_j(v,t)(1+\varepsilon_j^w(v,t))]^{\psi}}{[s_j(v,t)(1+\varepsilon_j^s(v,t))]^{\psi}} \tag{C.27}$$

$$l_j(v,t) = \frac{\bar{A}\hat{Y}_j(v,t)}{\hat{Q}_j(v,t)}\left[(1-B) + B[\hat{Q}_j(v,t)^{-\alpha}\chi_j(v,t)]^{\frac{\psi-1}{\psi}}\right]^{\frac{-\psi}{\psi-1}}, \tag{C.28}$$

where $m_j(v,t)$ is immediate from $\chi_j(v,t)$ and $\varepsilon_j^w(v,t)$ and $\varepsilon_j^s(v,t)$ are the inverse labor supply elasticities for w and s, respectively.

Let $\pi_j(v,t) = \frac{\Pi_j(v,t)}{Y(t)}$ be the profit per unit of output. Combining (C.21) with the firm's labor

requirement functions yields

$$\pi_j(v,t) = \max_{\hat{P}} \left[\hat{P}^{1-\sigma} - \left\{ w \left(\frac{\bar{A}}{\hat{Q}_j(v,t)} \hat{P}^{-\sigma} \right) + \hat{Q}_j(v,t)^{\alpha} s \left(\frac{\bar{A}}{B \hat{Q}_j(v,t)^{1-\alpha}} \hat{P}^{-\sigma} \right) \right\} \frac{\bar{A}}{\hat{Q}_j(v,t)} \hat{P}^{-\sigma} \right] \quad \text{(C.29)}$$

We can rewrite the FOC for the optimal price in terms of elasticities, yielding

$$\hat{P}(\hat{Q}) = \underbrace{\frac{\sigma}{\sigma - 1}}_{\text{CES markup}} \frac{\bar{A}}{\hat{Q}_j(v,t)} \underbrace{\left[(1-B)^{\psi} [w_j(v,t)(1+\varepsilon_j^w(v,t))]^{1-\psi} + B^{\psi} \hat{Q}_j(v,t)^{\alpha(1-\psi)} [s_j(v,t)(1+\varepsilon_j^s(v,t))]^{1-\psi} \right]^{\frac{1}{1-\psi}}}_{\text{monopsonist's marginal cost}} \quad \text{(C.3}$$

where $w(l)$ and $s(m)$ are the inverse labor supply curves in (C.16), and

$$\varepsilon_j^w(v,t) \equiv \frac{\partial w(l(\hat{Q}_j(v,t)))}{\partial l(\hat{Q}_j(v,t))} \frac{l(\hat{Q}_j(v,t))}{w(l(\hat{Q}_j(v,t)))}, \; \varepsilon_j^s(\hat{Q}_j(v,t);s) \equiv \frac{\partial s(m(\hat{Q}_j(v,t)))}{\partial m(\hat{Q}_j(v,t))} \frac{m(\hat{Q}_j(v,t))}{s(m(\hat{Q}_j(v,t)))} \quad \text{(C.31)}$$

are the elasticity of the inverse labor supply curve.[3]

While in the prior section we derived the probability that a worker chose a job, in order to calculate the elasticity of labor supply we need to compute the expected number of efficiency units that a firm would get given a posted wage w. Workers leave their outside option if the take home pay at their best firm is greater than $R^{1+\frac{1}{\varphi}}$. We can derive the expected number of efficiency units for job k by integrating

$$N_k = \frac{\tilde{\theta}_\tau S_k}{w_k} \int_{R_\tau^\beta}^{\infty} y^{-\tilde{\theta}_\tau} \, e^{-\Lambda_\tau y^{-\tilde{\theta}_\tau}} \, dy \quad \text{(C.32)}$$

This gives us an expression involving the lower incomplete Gamma function, which appears because of the truncation of the domain

$$N_k = \frac{S_k}{w_k} \Lambda_\tau^{1/\tilde{\theta}_\tau - 1} \, \gamma\Big(1 - \tfrac{1}{\tilde{\theta}_\tau}, \, \Lambda_\tau R_\tau^{-\theta_\tau}\Big). \quad \text{(C.33)}$$

This gives us an expression for aggregate labor supply as a function of wages,

$$l(w) = \omega K \frac{S_k}{w_k} \Lambda_l^{1/\tilde{\theta}_l - 1} \, \gamma\Big(1 - \tfrac{1}{\tilde{\theta}_l}, \, \Lambda_l R_l^{-\theta_l}\Big) \quad \text{(C.34)}$$

where the expression for type m is analogous.

We can take logs and differentiate this with respect to w to compute the elasticities:

[3] In this expression, the $\bar{A}$ term appears since wages are scaled by aggregate output, while the $\hat{Q}$ reflects the impact of relative productivity differences on unit labor costs across firms with the same industry.

$$\varepsilon_j^\tau = (\theta_\tau - 1) + \frac{T_\tau \theta_\tau w_k^{\theta_\tau} \mathbb{E}_W \left[\Lambda_\tau^{\frac{1}{\theta_\tau}-2} \left(\left(\frac{1}{\theta_\tau} - 1\right) \gamma(1 - \frac{1}{\theta_\tau}, \Lambda_\tau R_\tau^{-\theta_\tau}) + \left(\Lambda_\tau R_\tau^{-\theta_\tau}\right)^{1-\frac{1}{\theta_\tau}} e^{-\Lambda_\tau R_\tau^{-\theta_\tau}} \right) \right]}{\mathbb{E}_W \left[\Lambda_\tau^{\frac{1}{\theta_\tau}-1} \gamma(1 - \frac{1}{\theta_\tau}, \Lambda_\tau R_\tau^{-\theta_\tau}) \right]} \tag{C.35}$$

Finally, we need to pin down $\bar{A}$. By combining the production function with the CES demand function, we know that

$$\frac{Y_j(v,t)}{\bar{Q}_j(t)} = \hat{Q}_j(v,t) l_j(v,t) = \frac{Y(t)}{P_j(t)\bar{Q}_j(t)} \hat{P}_j(v,t)^{-\sigma} = \frac{\bar{A}\bar{Q}(t)}{\bar{Q}_j(t)} \frac{\bar{Q}_j(t)}{\bar{Q}(t)} \hat{P}_j(v,t)^{-\sigma} = \bar{A}\hat{P}_j(v,t)^{-\sigma}, \tag{C.36}$$

which we can rearrange and integrate to get

$$\bar{A} = \frac{\int_0^1 \int_0^1 \hat{Y}_j(v,t) \mathrm{d}j \, \mathrm{d}v}{\int_0^1 \int_0^1 \hat{P}_j(v,t)^{-\sigma} \mathrm{d}j \, \mathrm{d}v} = \int_0^1 \int_0^1 \hat{Q}_j(v,t) \left[(1-B) l_j(v,t)^{\frac{\psi-1}{\psi}} + B[\hat{Q}_j(v,t)^{-\alpha} m_j(v,t)]^{\frac{\psi-1}{\psi}} \right]^{\frac{\psi}{\psi-1}} \mathrm{d}j \, \mathrm{d}v, \tag{C.37}$$

where the last equality follows because $\int_0^1 \hat{P}_j(v,t)^{-\sigma} \mathrm{d}v = \int_0^1 \frac{Y_j(v,t)}{Y_j(t)} \mathrm{d}v = 1$. The last equality says that we can also write (C.37) as a labor-weighted average of relative labor productivity $\hat{Q}_j(v,t)$ times the aggregate number of efficiency units of labor.

C.2 Dynamic model

Next, we discuss a number of technical details about the dynamic part of our model.

Illustration of the Business-Time Construction

Appendix Figure A.4 illustrates how the subordinator $\Gamma_j(t)$ maps calendar time into business time over two years for a single firm. Panel A plots the stochastic evolution of $\Gamma_j(t)$: business time advances only when $\Gamma_j(t)$ jumps, so periods of rapid industry innovation correspond to steeper segments of the business-time profile. Panel B shows the relative productivity $\hat{Q}_j(v,t)$ as a function of business time: it drifts downward at the deterministic rate $\tilde{g}$ (reflecting ongoing improvements in average industry productivity) and occasionally jumps upward when the firm successfully innovates—resembling a standard Poisson process but in business time. Panel C shows the same process in calendar time: whenever $\Gamma_j(t)$ jumps (an industry innovation burst), $\hat{Q}_j(v,t)$ shifts downward; when the firm itself innovates, it jumps upward. This construction introduces industry-level stochastic shocks that would be absent from a standard growth framework with deterministic industry dynamics.

External innovation: technical details. For each product line in which the firm is the leading producer, external business-stealing improvements arrive according to

$$N_j^x(v,t) \sim \text{Poisson}\big(x_j(v,t^-)\,\mathrm{d}\Gamma_j(t)\big)\,. \tag{C.38}$$

These improvements are undirected: when $N_j^x(v,t) > 0$, the firm draws $N_j^x(v,t)$ varieties at random and improves their productivity by λ. A fraction $1 - p_r$ of these innovations reset productivity in the affected lines to $\lambda\,\bar{Q}_j(t)$, ensuring stationarity of the distribution of relative productivities $\hat{Q}_j(v,t)$. When business time advances by more than an infinitesimal amount (i.e., $\mathrm{d}\Gamma_j(t) > 0$), it is possible—though with very small probability—for a variety to experience both an internal and an external innovation in the same instant. For tractability, we assume that if this occurs, the gains from internal innovation accrue to the firm that ultimately acquires the product line.

Innovation by new entrants: technical details. Let $e_j(t)$ denote the mass of new entrants. The number of external innovations arriving at a given product line from entrants is

$$N_j^e(v,t) \sim \text{Poisson}\big(e_j(t^-)\,\mathrm{d}\Gamma_j(t)\big)\,. \tag{C.39}$$

Because innovation occurs in bursts, multiple firms may innovate at the same instant within the same product line. In this unlikely event, we randomly select one of the successful external innovators to become the new leading producer—capturing the idea that external innovations occur sequentially in business time, so the winner is the firm that attains the latest improvement during a burst.

Worker productivity dynamics. To describe how match quality evolves, we work with a transformed index

$$z^*(i,k,t) \equiv \Phi^{-1}\Big[F_{\tau(i)}\Big(z_{j(i,k,t)}(i,v(i,k,t),t)\Big)\Big]\,, \tag{C.40}$$

which maps the Fréchet productivity draw z into a standard normal variable z^*. This inverse-transform representation allows us to specify simple OU dynamics for z^* while preserving a Fréchet distribution for actual productivity at all times. When a product line is displaced by creative destruction, we reset only that job's transformed match quality to $z^*(i,k,t) \sim N(0,1)$ and assign new industry and variety indices, drawn from the product lines that changed their active producer during the innovation burst. This approach preserves the Fréchet distribution of match-specific productivity.

Dynamics of $\hat{q}$, $\bar{q}$, and aggregate growth

Let $q_j(v,t) \equiv \log Q_j(v,t)$. Under the assumptions in the text, we can write the dynamics of $q_j(v,t)$ as follows

$$\mathrm{d}q_j(v,t) \;:=\; \log\lambda\,\mathrm{d}N_j(v,t) - \underbrace{\big(1-\tilde{D}_j(v,t)\big)\hat{q}_j(v,t^-)}_{\text{reset to } \bar{q}_j(t)}\,, \tag{C.41}$$

where, given that there are n jumps at the product line level, the probability that the product line is not 'reset' equals

$$\Pr(\tilde{D}_j(v,t)=1|dN_j(v,t)=n) \;=\; \left[\frac{I[\hat{q}_j(v,t^-)]+\tau(1-p_r)}{I[\hat{q}_j(v,t^-)]+\tau}\right]^n \equiv \gamma_n(\hat{q}_j(v,t^-)). \tag{C.42}$$

Hence, we can calculate (using $\hat{q}$ as shorthand for $\hat{q}_j(v,t^-)$)

$$E\left[\frac{\mathrm{d}Q_j(v,t)}{Q_j(v,t^-)}|d\Gamma_{j,t}=\mu\right] = e^{-(I(\hat{q})+\tau)\mu}\sum_{k=1}^{\infty}\frac{\left([I(\hat{q})+\tau]\,\mu\right)^k}{k!}\left[\gamma_k(\hat{q})(\lambda^k-1)+(1-\gamma_k(\hat{q}))\left(\frac{\lambda^k}{\hat{Q}_j(v,t^-)}-1\right)\right]. \tag{C.43}$$

Let's next explore the dynamics of $\bar{Q}_j(t)\equiv\int_0^1 Q_j(v,t)dv$, where $\bar{q}_j(t)\equiv\log\bar{Q}_j(t)$. We can write the aggregate growth rate of industry productivity as

$$\frac{\mathrm{d}\bar{Q}_j(t)}{\bar{Q}_j(t^-)}=\int_0^1\hat{Q}_j(v,t)\cdot\frac{\mathrm{d}Q_j(v,t)}{Q_j(v,t^-)}\mathrm{d}v = E\left[\hat{Q}_j(v,t)\frac{\mathrm{d}Q_j(v,t)}{Q_j(v,t^-)}\right] = \mathbf{1}(\mathrm{d}\Gamma_{jt}=\mu)E\left[\hat{Q}_j(v,t)\frac{\mathrm{d}Q_j(v,t)}{Q_j(v,t^-)}|\Gamma_{jt}=\mu\right].$$

So it follows that

$$E\left[\hat{Q}_j(v,t)\tfrac{\mathrm{d}Q_j(v,t)}{Q_j(v,t^-)}|d\Gamma_{j,t}=\mu\right] = E\left\{e^{-(I(\hat{q})+\tau)\mu}\sum_{k=1}^{\infty}\frac{\left([I(\hat{q})+\tau]\,\mu\right)^k}{k!}\left[\gamma_k(\hat{q})(\lambda^k-1)\hat{Q}_j(v,t^-)+(1-\gamma_k(\hat{q}))\left(\lambda^k-\hat{Q}_j(v,t^-)\right)\right]\right\}.$$

$$\text{(if } I(\hat{q})=I) = e^{-(I+\tau)\mu}\sum_{k=1}^{\infty}\frac{\left([I+\tau]\,\mu\right)^k}{k!}\left[\gamma_k(0)(\lambda^k-1)\underbrace{E\{\hat{Q}_j(v,t^-)\}}_{=1}+(1-\gamma_k(0))\left(\lambda^k-E\{\hat{Q}_j(v,t^-)\}\right)\right]$$

$$= e^{-(I+\tau)\mu}\sum_{k=1}^{\infty}\frac{\left([I+\tau]\,\mu\right)^k}{k!}(\lambda^k-1) = \exp((\lambda-1)(I+\tau)\mu)-1.$$

where the second line replaces $\gamma_k(\hat{q})$ with $\gamma_k(0)$ to reflect the fact that $\gamma_k(\hat{q})$ is constant when I is constant, which allows us to take the expectation inside the sum.

If we instead work in logs, (C.2) implies that if I is constant (an assumption which will hold in our quantitative model given restrictions on $\Phi(\hat{Q})$ that we will describe below), then

$$\mathrm{d}\bar{q}_j(t)=(\lambda-1)(I+\tau)\mathrm{d}\Gamma_{j,t}\equiv\tilde{g}\,\mathrm{d}\Gamma_{j,t}. \tag{C.44}$$

When I is not constant, we need to instead define $\tilde{g}$ using the more complicated expression in the top line of (C.2). Finally, the aggregate growth rate equals

$$g\mathrm{d}t=\frac{\mathrm{d}\bar{Q}(t)}{\bar{Q}(t)}=E\left[\frac{\bar{Q}_j(t^-)}{\bar{Q}(t)}\frac{\mathrm{d}\bar{Q}_j(t)}{\bar{Q}_j(t^-)}\right]=\underbrace{E\left[\frac{\bar{Q}_j(t^-)}{\bar{Q}(t)}\right]}_{=1}E\left[\frac{\mathrm{d}\bar{Q}_j(t)}{\bar{Q}_j(t^-)}\right]=\frac{1}{\mu}[\exp((\tilde{g}-\delta)\mu)-1]\mathrm{d}t\approx(\tilde{g}-\delta)\mathrm{d}t, \tag{C.45}$$

where we used the fact that the increments to industry productivity are iid across industries to go from the second to the third equality. While $g\neq\tilde{g}$ due to a slight adjustment for Jensen's inequality, the difference between the two is small when μ isn't large.

Approximating the stationary distribution of $\hat{q}$

Next, we describe how we approximate the stationary distribution of $\hat{q}_j(v,t)$ numerically. Combining (C.41) with (C.44) yields the following law of motion for the $\hat{q}_j(v,t)$ of the most efficient producer in each industry and variety:

$$\mathrm{d}\hat{q}_j(v,t) = \log\lambda\,\mathrm{d}N_j(v,t) - (1-\tilde{D}_j(v,t))\hat{q}_j(v,t^-) - \tilde{g}\,\mathrm{d}\Gamma_{j,t}. \tag{C.46}$$

For brevity, we omit j and v subscripts, since $\hat{q}$ is iid across varieties and industries along the BGP.

HJB, Innovation costs $\Phi(\hat{q})$, and ODEs for product line valuations

Economic interpretation of the HJB. The Hamilton–Jacobi–Bellman equation (39) in the main text trades off four forces. First, current flow profits $\pi(\hat{q})$ reflect the firm's static markup earnings on its existing product line. Second, the investment cost terms c_I and c_x represent the resources devoted to internal and external innovation. Third, the term $x\,\mathcal{G}_x V$ captures the expected gain from a successful business-stealing innovation: with intensity x, the firm leapfrogs a competitor and adds a new product line to its portfolio. Fourth, the final bracketed term governs how internal innovation and creative destruction jointly shape the value of the existing product line through changes in relative productivity: internal innovation raises $\hat{q}$ by discrete steps of size $\log\lambda$, industry-level progress erodes relative productivity at rate $\tilde{g}\,d\Gamma_j$, and with intensity $\bar{x}+\bar{e}$ the product line is lost to creative destruction—reducing the continuation value to zero. The required return on the left-hand side equals $r-g$ because values are expressed relative to aggregate output.

Next, we want to compute the present discounted value of profits. Assuming the two-stage game equilibrium concept (which eliminates limit pricing), then profits depend only on the firm's $\hat{q}$ and we can drop κ from the state space. Suppose the flow payoff is $Y(t)\pi(\hat{q}_j(v,t))$, discounted at rate r, where $\pi(0)=0$. Supposing $\frac{\dot{Y}(t)}{Y(t)} = g\,\mathrm{d}t$, then we can define the value of a product line per unit aggregate output as

$$V(\hat{q}_j(v,t)) = \mathbb{E}\Big[\int_0^\infty e^{-(r-g)\,t}\left[\pi(\hat{q}) - c_I\left(I;\hat{q}\right) - c_x\left(x;\hat{q}\right)\right]\mathrm{d}t\Big], \tag{C.47}$$

where profits and costs include additional product lines which obtain as a result of external innovation. We can write this valuation recursively as

$$\begin{aligned}(r-g)\,V(\hat{q}) \;=\; \max_{I(\hat{q}),x}\;& \pi(\hat{q}) - c(I,x;\hat{q}) + x\mathcal{G}_x V \\ &+\frac{1}{\mu}\Big[e^{-(I(\hat{q})+\tau)\mu}\sum_{k=0}^{\infty}\frac{\left(I(\hat{q})\,\mu\right)^k}{k!}\;V(\hat{q}-\tilde{g}\,\mu+k\log\lambda) - V(\hat{q})\Big],\end{aligned} \tag{C.48}$$

where the second line – which captures the benefits of internal innovation and loss of product lines

via displacement – obtains from the following direct calculation

$$\begin{aligned}
[E_t V(\hat{q}_t) | d\Gamma_t = \mu, N_t^x = 0] &= e^{-(I(\hat{q}_{t^-})+\tau)\mu}\, V(\hat{q}_{t^-} \underbrace{-\tilde{g}\,\mu}_{\text{competitor price impact}}) \\
&+ \sum_{k=1}^{\infty} e^{-(I(\hat{q}_{t^-})+\tau)\mu} \frac{\left((I(\hat{q}_{t^-})+\tau)\,\mu\right)^k}{k!} \underbrace{\left[\left[\frac{I(\hat{q}_{t^-})}{I(\hat{q}_{t^-})+\tau}\right]^k V(\hat{q}_{t^-} - \tilde{g}\,\mu + k\log\lambda)\right]}_{\text{improve product line}} \\
&+ \sum_{k=1}^{\infty} e^{-(I(\hat{q}_{t^-})+\tau)\mu} \frac{\left((I(\hat{q}_{t^-})+\tau)\,\mu\right)^k}{k!} \underbrace{\left[(1 - \left[\frac{I(\hat{q}_{t^-})}{I(\hat{q}_{t^-})+\tau}\right]^k)\, 0\right]}_{\text{lose product line}} \\
&= e^{-(I(\hat{q}_{t^-})+\tau)\mu} \sum_{k=0}^{\infty} \frac{\left(I(\hat{q}_{t^-})\,\mu\right)^k}{k!} V(\hat{q}_{t^-} - \tilde{g}\,\mu + k\log\lambda). \qquad \text{(C.49)}
\end{aligned}$$

Notice that increases in τ decrease the continuation value through a higher likelihood that a competitor steals the line, which drops the value of profits from the existing line to zero.

Next, we need to work out what is $\mathcal{G}_x V$. By direct calculation and an application of the tie-breaking rules for multiple jumps discussed in the main text, we get

$$\mathcal{G}_x V = \mathbb{E}\left\{ \frac{\exp(-[I(\hat{q})+\tau]\mu)}{P_{x\geq 1}} \sum_{k=1}^{\infty} \frac{[I(\hat{q})+\tau]^k \mu^k}{k!} \left[\underbrace{\sum_{n=1}^{k} \sum_{j=1}^{n} \frac{k!\, p_x(\hat{q})^j p_e(\hat{q})^{n-j} p_I(\hat{q})^{k-n}}{j!(n-j)!(k-n)!}}_{\substack{\text{Prob of } n \text{ bus stealing jumps} \\ \text{and at least 1 } x \text{ jump}}} \frac{1}{n} \left[\begin{array}{l} p_r^n V(k\log\lambda + \hat{q} - \tilde{g}\mu) \\ +(1-p_r^n) V(k\log\lambda - \tilde{g}\mu) \end{array} \right] \right] \right\}, \qquad \text{(C.50)}$$

where $P_{x\geq 1} = 1 - \exp(-x\mu)$ and

$$p_I(\hat{q}) = \frac{I(\hat{q})}{I(\hat{q})+\tau}, \qquad p_x(\hat{q}) = \frac{x}{I(\hat{q})+\tau}, \qquad p_e(\hat{q}) = \frac{e}{I(\hat{q})+\tau}. \qquad \text{(C.51)}$$

The $1/n$ factor reflects what happens if two firms have stealing innovations at the same time. This captures the probability that a firm has the last of the innovations.

Analogously, we can compute $\mathcal{G}_e V$, the same object for an entry-related jump.

$$\mathcal{G}_e V = \mathbb{E}\left\{ \frac{\exp(-[I(\hat{q})+\tau]\mu)}{P_{e\geq 1}} \sum_{k=1}^{\infty} \frac{[I(\hat{q})+\tau]^k \mu^k}{k!} \left[\underbrace{\sum_{n=1}^{k} \sum_{j=1}^{n} \frac{k!\, p_e(\hat{q})^j p_x(\hat{q})^{n-j} p_I(\hat{q})^{k-n}}{j!(n-j)!(k-n)!}}_{\substack{\text{Prob of } n \text{ bus stealing jumps} \\ \text{and at least 1 } e \text{ jump}}} \frac{1}{n} \left[\begin{array}{l} p_r^n V(k\log\lambda + \hat{q} - \tilde{g}\mu) \\ +(1-p_r^n) V(k\log\lambda - \tilde{g}\mu) \end{array} \right] \right] \right\}, \qquad \text{(C.52)}$$

where $P_{e\geq 1} = 1 - \exp(-e\mu)$.

In order to characterize the size distribution of the number of product lines, it's also important to know the probability that a firm "wins" the product line conditional on an x jump (i.e., a successful

external innovation), which equals

$$P_{x \text{ success}} = \frac{\exp(-\tau\mu)}{1-\exp(-x\mu)} \sum_{k=1}^{\infty} \frac{(\tau\mu)^k}{k!} \frac{1}{k} \underbrace{\left[1-\left(\frac{e}{\tau}\right)^k\right]}_{\text{Prob of at least 1 x jump}}, \tag{C.53}$$

along with its counterpart for an external innovation by a new entrant, which is

$$P_{e \text{ success}} = \frac{\exp(-\tau\mu)}{1-\exp(-e\mu)} \sum_{k=1}^{\infty} \frac{(\tau\mu)^k}{k!} \frac{1}{k} \left[1-\left(\frac{x}{\tau}\right)^k\right]. \tag{C.54}$$

In practice, the probability of multiple jumps is very low in our calibrated model, so $P_{x \text{ success}} \approx 1$ and these additional adjustments are very minor.

As is discussed in the main text, we choose $\Phi(\hat{q})$ so that innovation intensities stay constant regardless of $\hat{q}$, which is conceptually similar to an assumption made by Peters (2020). He assumes that the costs of incremental innovation scale proportionally with the benefits, yielding a constant innovation rate. Here, we will define $\Phi(\hat{q})$ implicitly, in such a way that we can solve for it directly as part of an ODE which pins down firm valuations. The first order condition with respect to $I(\hat{q})$ equals

$$\Phi(\hat{q}) \frac{\zeta I(\hat{q})^{\zeta-1}}{\varphi_I} = \mathcal{G}_I V, \tag{C.55}$$

where

$$\mathcal{G}_I V \equiv \underbrace{e^{-(I(\hat{q})+\tau)\mu} \left[\sum_{k=0}^{\infty} \left(\frac{[I(\hat{q})\mu]^k}{k!}\right) \left[V\left(\hat{q}-\tilde{g}\,\mu+(k+1)\log\lambda\right)-V\left(\hat{q}-\tilde{g}\,\mu+k\log\lambda\right)\right]\right]}_{\mathcal{G}_I V} \tag{C.56}$$

If we rearrange and impose that $I(\hat{q})$ is constant, it follows that the following choice of $\Phi(\hat{q})$ is consistent with a constant innovation arrival rate of I:

$$\Phi(\hat{q}) = \frac{\varphi_I}{\zeta\,(I^*)^{\zeta-1}}\, e^{-(I+\tau)\mu} \left[\sum_{k=0}^{\infty} \left(\frac{[I\mu]^k}{k!}\right) \left[V\left(\hat{q}-\tilde{g}\,\mu+(k+1)\log\lambda\right)-V\left(\hat{q}-\tilde{g}\,\mu+k\log\lambda\right)\right]\right] \tag{C.57}$$

For the calibration in the paper, the implied function $\Phi(\hat{q})$ is plotted in Figure A.7 together with the ergodic distribution of $\hat{q}$.

Analogously, we can solve for φ_x which is consistent with a particular value of radical innovation by incumbents, x, which yields:

$$\varphi_x = \frac{\zeta\, x^{\zeta-1}}{\mathcal{G}_x\, V}. \tag{C.58}$$

If we assume that $\Phi(\hat{q})$ satisfies C.57 and plug in the optimal I and x into the HJB (C.48), we get the following ODE for firm valuations in terms of the optimal innovation arrival rates,

$$(r-g)\,V(\hat{q}) \quad = \quad \pi(\hat{q}) + x\left(1-\frac{1}{\zeta}\right)\mathcal{G}_x\, V$$

$$- \quad \frac{I}{\zeta}\, e^{-(I+\tau)\mu} \left[\sum_{k=0}^{\infty} \left(\frac{[I\mu]^k}{k!} \right) \left[V\left(\hat{q} - \tilde{g}\,\mu + (k+1)\log\lambda\right) - V\left(\hat{q} - \tilde{g}\,\mu + k\log\lambda\right) \right] \right] \tag{C.59}$$

$$+ \quad \frac{1}{\mu} \left[e^{-(I+\tau)\mu} \sum_{k=0}^{\infty} \frac{\left(I\,\mu\right)^k}{k!}\; V\left(\hat{q} - \tilde{g}\,\mu + k\log\lambda\right) - V(\hat{q}) \right],$$

which, once we've discretized over a $\hat{q}$ grid, reduces to a linear system of equations.

Stationary distribution of number of product lines: $\mu \downarrow 0$ limit

Here, we solve for the measure of product lines that are associated with firms with different numbers of products. If we take the limit as $\mu \downarrow 0$, we can solve the KFE for this distribution in closed form. Since the probability of moving down in the distribution of number of lines exceeds the probability of moving up, so there will be a well-defined size distribution. Here are the key accounting identities which capture h_n the mass of *product lines* at firms which currently have n lines:

$$\dot{h}_1 = z - (x+\tau)h_1 + \tau h_2 = 0 \tag{C.60}$$

$$\dot{h}_n = x h_{n-1} - (x+\tau)h_n + \tau h_{n+1} = 0. \tag{C.61}$$

In a BGP in which the mass of firms is constant, entry must offset exit, which implies that $z = \tau h_1$, so $h_1 = \frac{z}{\tau} = \frac{\tau - x}{\tau}$. Next, let's guess and verify that $h_n = c h_{n-1}$ satisfies the recursion. Let's start with the first equation. If we know that $h_1 = \frac{\tau - x}{\tau}$, then we get $x h_1 = \tau h_2$, which implies that $c = \frac{x}{\tau} = \frac{x}{x+z}$. Next, we verify that $c = \frac{x}{\tau}$ satisfies the general second order difference equation in (C.61) for any $n \geq 2$:

$$x h_{n-1} - c(x+\tau)h_{n-1} + \tau c^2 h_{n-1} = h_{n-1}[x - \frac{(x+\tau)x}{\tau} + \frac{x^2}{\tau}] = h_{n-1} \cdot 0. \tag{C.62}$$

It is also easy to verify that the pmf $h_n = \frac{\tau - x}{\tau} \left[\frac{x}{\tau}\right]^{n-1}$ integrates to one.

Since all firms have at least 1 product, the measure of total firms will be less than the unit measure of total products; instead, it will equal $1 \;/\; \mathbb{E}[N_{f,t} | N_{f,t} > 0]$. Hence, if we want to compute the share of *firms* of each size, which we will denote by p_n, we need to do the following transformation:

$$p_n = \frac{\frac{1}{n} h_n}{\sum_k \frac{1}{k} h_k} = \frac{-x}{\tau \log(1 - \frac{x}{\tau})} \left[\frac{x}{\tau}\right]^{n-1} \frac{1}{n}. \tag{C.63}$$

Stationary distribution of number of product lines: Exact calculation with finite μ

Here, we instead derive the distribution of the number of product lines without taking the limit as $\mu \downarrow 0$. To do this, analogously with Klette and Kortum (2004), we will work with probability generating functions (PGFs). Suppose that we have a firm which currently has $N_{f,t}$ product lines.

Then, the new number of product lines can be decomposed as

$$N_{f,t} = \begin{cases} \sum_{j=1}^{N_{f,t^-}} (S_{j,f,t} + I^x_{j,f,t}) & \text{if } \sum_{j=1}^{N_{f,t^-}} (S_{j,f,t} + I^x_{j,f,t}) > 0 \\ 1 & \text{otherwise} \end{cases}, \tag{C.64}$$

where $S_{j,f,t}$ captures product line survival and $I^x_{j,f,t}$ captures the successful arrival of new product lines. Let's define by $G(s) \equiv (p_\tau + (1-p_\tau)s)\exp(x\mu P_{x\text{ success}}(s-1))$ as the PGF of the change associated with an individual product line, where $p_\tau \equiv 1-\exp(-\tau\mu)$ is the probability of displacement during a single industry innovation burst. The $P_{x\text{ success}}$ term is necessary to capture the fact that there is a very small probability that a successful external innovation may not result in winning the market.

Using properties of PGFs, we get that

$$\mathbb{E}[s^{N_{f,t}} | N_{f,t^-} = n, \mathrm{d}\Gamma_{j(f),t} = \mu] = G(s)^n + G(0)^n(s-1), \tag{C.65}$$

where the last term reflects that, in the stationary distribution, exiting firms are immediately replaced by new entrants with 1 product line. Now, let's find $H(s)$, the PGF of the stationary distribution.

$$\begin{aligned} G(s) \equiv \mathbb{E}[s^{N_{f,t}}] &= \mathbb{E}[E(s^{N_{f,t}} | N_{f,t^-}, \mathrm{d}\Gamma_{j(f),t} = \mu)] \qquad \text{(law of iterated expectations)} & \text{(C.66)} \\ &= \mathbb{E}[G(s)^{N_{f,t^-}}] + \mathbb{E}[G(0)^{N_{f,t^-}}](s-1) = H(G(s)) + H(G(0))(s-1). & \text{(C.67)} \end{aligned}$$

Let's denote by $E \equiv H(G(0))$, which captures the mass of entering firms. Let's next characterize $H(s)$ by recursive substitution:

$$\begin{aligned} H(s) &= H(G(s)) + E(s-1) & \text{(C.68)} \\ H(G(s)) &= H(G^{\circ 2}(s)) + E(G(s)-1) & \text{(C.69)} \\ H(s) &= 1 + E\sum_{j=0}^{\infty}(G^{\circ j}(s) - 1), & \text{(C.70)} \end{aligned}$$

where here we use the fact that $\lim_{N\to\infty} H(G^{\circ N}(s)) = 1$ for all $s \in [0,1]$, which will hold as long as all firms die out in the long run (ie each firm's number of products shrinks over time in expectation).

While $H(\cdot)$ does not satisfy a recursion which we can evaluate in closed form, we can use Fourier inversion techniques to recover the pmf of the stationary distribution. In particular, an application of Cauchy's integral formula implies that for an analytic function

$$p_n = \frac{H^{(n)}(0)}{n!} = \frac{1}{2\pi i}\oint_{|s|=r} \frac{H(s)}{s^{n+1}}\, ds, \tag{C.71}$$

where we take an integral around any circle $|s| = r \leq 1$. Next, if we write points on the unit circle

as $s = \exp(i\theta)$, for $\theta \in [0, 2\pi]$, with $ds = i\exp(i\theta)d\theta$

$$p_n \;=\; \frac{1}{2\pi i}\int_0^{2\pi} H(e^{i\theta})\, e^{-i(n+1)\theta}\,[i\exp(i\theta)d\theta] \;=\; \frac{1}{2\pi}\int_0^{2\pi} H(e^{i\theta})\, e^{-in\theta}\, d\theta. \tag{C.72}$$

Now, if we approximate the integral using the trapezoid rule, we get the following discrete Fourier transform formula:

$$p_n \;=\; \frac{1}{M}\sum_{k=0}^{M-1} H(e^{2\pi ik/M})\, e^{-2\pi ikn/M}. \tag{C.73}$$

Next, we need to pin down what E is. If we want the measure of firms to stay constant along the BGP, then entry needs to offset exit.

What is the rate of new entrants, E? Recall that mass e of entrants enter with probability $d\Gamma_{j,t}$. The probability they are successful is given by $P_{e\text{ success}}$ in equation (C.54) above, which will typically be very close to 1 since the probability of multiple successful business stealing innovations in the same $\mathrm{d}\Gamma_{j,t}$ burst is very unlikely. Hence, the total mass of successful entrants will be $E = (e\mu)P_{e\text{ success}}$.

If we define $E = (e\mu)P_{e\text{ success}}$, we solve for the measure of firms which have each number of product lines. This will integrate to a number equal to $1\ /\ E[N_{f,t}|N_{f,t} > 0]$. An alternative way to pin down E is to solve for the level of E which sets $H(0) = 0$. This works because $H(0)$ is equal to p_0, the mass of firms with zero product lines. Therefore, setting E so that $H(0) = 0$ guarantees that each dead firm is replaced by a newborn one and that p_n captures the conditional probability distribution of firm sizes for all firms with positive numbers of product lines. Reassuringly, both methods yield the same conditional pmf for positive product lines, which closely coincides with what obtains from the $\mu \downarrow 0$ limit. Figure 16 illustrates this latter point in a numerical example.

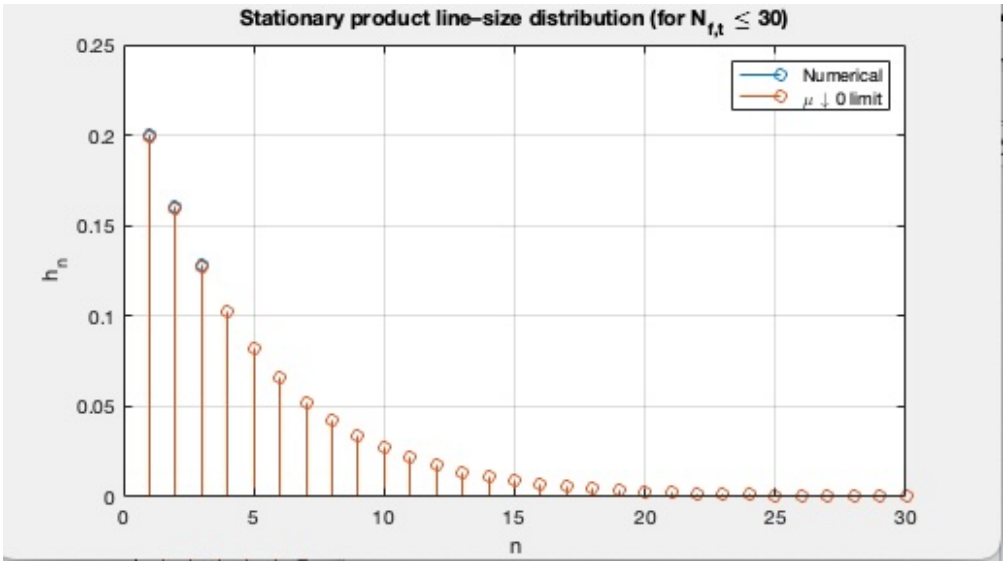


Figure 16: Comparison of h_n in exact calculation vs $\mu \downarrow 0$ analytical solution

D Numerical Implementation

D.1 Calibration Parameters

Model Mapping to Observable Measures of Innovation

The key explanatory variables in Section 2 consist of the value of innovation by the firm and its competitors. To interpret these objects through the lens of the model, consider the change in the market value of a product line as a result of innovation,

$$\Delta V(q_{f,l})(v,t^+) = \max\left(V\left(\hat{q}_{f,l}(v,t^+)\right) - V\left(\hat{q}_{f,l}(v,t^-)\right), 0\right). \tag{C.74}$$

This value can increase for two reasons: the firm has successfully innovated internally (raising the productivity advantage and hence the value of that product line), or it has innovated externally (adding a new product line to its portfolio).

The KPSS value of innovation for firm f in period t is the sum of this increase along the firm's product lines,

$$KPSS_{f,t} = \sum_{n=1}^{N_{f,t}} \Delta V(q_{f,l})(v,t^+), \tag{C.75}$$

which maps directly into the numerator of our empirical measure of firm innovation (10).

Innovation by competing firms is defined analogously as the sum of KPSS values across all other firms in industry j:

$$KPSS_j = \int_0^1 \Delta KPSS_{f,t}\, df \propto \mathrm{d}\Gamma_{j,t}. \tag{C.76}$$

Given the continuum of firms within each industry, the industry-level value of KPSS is simply proportional to the realization of the industry innovation process defined in equation (27). This proportionality means that our empirical regressions of worker earnings on competitor innovation identify the causal effect of the industry-level subordinator.

Parameter Details

Here we provide details on the numerical implementation of the indirect inference procedure discussed in Section 3.4. To enforce linear constraints that guarantee the existence of the expectation of the Freéchet distribution, we parameterize θ_τ in terms of logs, and constrain $\theta_m \geq \theta_l$ for computational efficiency for the optimizer. We choose an outside option, R, so that 2% of type l workers are self-employed; because of greater skill dispersion for managers this implies substantially higher rates of self-employment for type m.

To find the conditional growth rate we target an aggregate growth rate of 2%, which implies a combination of a conditional growth rate $\tilde{g}$ and a depreciation parameter, δ. Given a $\tilde{g}$ or a δ, we

have the other, given by the following expression

$$\delta = \tilde{g} - \frac{\log(\mu g + 1)}{\mu}$$

D.2 Finding the stationary distribution

We approximate the product-line level quality on $\log\left(\hat{Q}\right)$ on an equally spaced grid with 201 points that has maximum of 2.3 and a minimum of -4.6. Let $\hat{Q}_j(v,t)$ be the quality of the product line and $\hat{q}_j(v,t)$ be the log-quality. Given a $d\Gamma$ shock, there are three possible things that can happen to a line: the firm can fail to innovate in that product line, successfully innovate in that product line (an I jump), or it can be displaced by a new firm.

Conditional on a $\mathrm{d}\Gamma_t$ shock, $\mathrm{d}\hat{q}_t$ has a finite support and probabilities given by the Poisson distribution. Specifically,

$$\mathrm{d}\hat{q}_t = \begin{cases} k\log\lambda - \tilde{g}\mathrm{d}\Gamma_t & \text{w/ prob. } \exp(-(I+\tau)\mathrm{d}\Gamma_t)\frac{[(I+\tau)\mathrm{d}\Gamma_t]^k}{k!}\left[\frac{I+\tau(1-p_r)}{I+\tau}\right]^k \\ k\log\lambda - \tilde{g}\mathrm{d}\Gamma_t - \hat{q}_{t^-} & \text{w/ prob. } \exp(-(I+\tau)\mathrm{d}\Gamma_t)\frac{[(I+\tau)\mathrm{d}\Gamma_t]^k}{k!}\left(1-\left[\frac{I+\tau(1-p_r)}{I+\tau}\right]^k\right) \end{cases}. \quad \text{(C.77)}$$

Notice that the process is constant whenever $\mathrm{d}\Gamma_t = 0$, so we can solve the stationary as if we are operating in discrete time, iterating on the stationary distribution across $\mathrm{d}\Gamma_t$ increments until convergence.

Since we operate on a discrete, bounded support, we impose clamps at the boundaries of the state space. In other words, if a jump moves $\hat{q}$ outside the state space, we replace the point in the support in (C.77) with the relevant boundary. Second, since the points in the support will typically not live exactly on the $\hat{q}$ over which we discretize, we construct a second stage lottery between the two adjacent grid points which has a mean equal to the desired point in the support. Since we work with a relatively fine grid in practice, this second stage lottery introduces very little additional randomness in the associated distribution. Figure A.7 plots our approximated distribution of $\hat{q}$.

Our grid is log-spaced while the quality per unit aggregate quality that is held constant along a balanced growth path is not in logs. Thus, when we approximate the ODEs along our grid, we linearly interpolate between adjacent points with weights that maintain the expected change in $\hat{Q}_t$.

In principle, there is no limit to the number of possible jumps that can occur for a single product line for a single firm over a single $d\Gamma$ shock because business time is no longer continuous but instead proceeds with a discrete step, μ. The probability of getting k jumps with rate I is

$$P(N=k) = \frac{(I\mu)^k}{k!}e^{-I\mu}$$

In practice, we choose a K_{max} such that $P(k > K_{max}) < \text{tol}$ where the tolerance is set to $1e-10$. This is not quite machine precision, but it allows us to work with a finite number of jumps and approximate them with a transition matrix along a discrete state-space.

To compute the stationary distribution, we solve the standard eigenvalue problem using Matlab's *eigs* command, normalizing the output to sum to one. We confirm the fixed point output by the solver by checking that

$$\pi_{\hat{q}} T_{\hat{q}} = \pi_{\hat{q}},$$

where $T_{\hat{q}}$ is the product line transition matrix and $\pi_{\hat{q}}$ is the steady state distribution of product line relative qualities. You can find the stationary distribution of $\hat{q}$ plotted in figure A.7.

D.3 Clearing Labor Markets

Recall that we do not have skill prices $w(\hat{Q})$ and $s(\hat{Q})$ in closed form, but instead have to solve for them by clearing markets. Given the distribution of $\hat{Q}$ which we find above, we can solve for wages that clear the supply and demand first order conditions discussed in the static equilibrium part of Section C.

We begin with two normalization restrictions. As long as T_τ for $\tau \in \{m, l\}$ is homogenous across jobs, we can't separately identify w from the number of efficiency units of labor. We choose $T_m = T_l$ so that the average number of efficiency units of each type equals one for type l, ie

$$T_l = \Gamma\left(1 - \tfrac{1}{\theta_l}\right)^{-1} \tag{C.78}$$

The functions $w(l)$ and $s(m)$ are only defined implicitly. To solve for them, we augment the system of equations to solve to include two labor market clearing conditions for each $\hat{Q}$ described above in equations C.27 and C.28. where $\mathcal{W}_i$ and $\mathcal{S}_i$ are defined as in the main text.

So, we can characterize the equilibrium in terms of the tuple $\{\mathbf{s}, \mathbf{w}, \hat{\mathbf{P}}, T_l, T_m, \bar{A}\}$ which satisfies profit maximization, labor supply functions (C.35), aggregate quality definitions(C.37), and satisfies the labor market clearing and normalization (C.78) equations for both workers and managers.

Given a guess of $\{\mathbf{w}, \mathbf{s}\}$, we can use the equations above to compute all other equilibrium objects. To compute the equilibrium conditions from a guess of $\{\mathbf{w}, \mathbf{s}\}$, we first compute the aggregate number of efficiency units (up to the scaling factors T_l and T_m) supplied of each type of labor.

This requires computing the expectation term which appears in (C.16), which we accomplish by simulating a large number of potential draws of peer firms' wages. We do this numerically by

simulating one thousand possible draws of outside options where workers each face K of them total. To accomplish this, in each state of the discrete state-space, we assume at least one job the worker faces is from that space, and then draw $K-1$ jobs from the multinomial distribution implied by the steady state distribution of product lines, $\pi_{\hat{q}}$. This lets us compute the distribution of outside options, that we can then plug into expressions C.16 and C.35. We need these for each type in order to compute the vector of labor supply elasticities.

This gives us the ingredients we need to guess and check with an optimizer to find the appropriate skill price functions. Within the objective function we use (C.78) and its analogue for managers to compute the scaling factors T_l and T_m, (C.37) to compute $\bar{A}$, (C.30) to compute prices and finally we use the difference between the labor supply and demand conditions to check whether the labor markets clear at the proposed wages.

In practice, we rely on Matlab's *lsqnonlin* function with an initial guess of constant wages, an objective tolerance of 1e-6, and a maximum number of function evaluations of $1e5$. If the first order conditions do not clear, we add a very large constant to the loss function so that the objective function is pushed away from those parameters.

D.4 Simulating Firms and Workers

Given the schedule of wages, the transition process for product-line level states, and the stationary distribution for initial conditions, we can then directly simulate and then compute our target moments that are analogous to the data. Before we discuss estimating those regression coefficients, we will discuss the simulation process for firms and workers.

We draw initial conditions for N_w workers from the stationary distribution of product-line qualities $\pi_{\hat{q}}$. Each worker has K product lines that they track over time–this is the set of jobs they see available at each instant. At each point in time, workers maximize their income and they can switch costlessly between these K jobs.

Rather than simulating on the discrete grid, we simulate directly from the ODEs. While time is continuous, in practice we must simulate over a discrete time-step Δ_t. For convenience (and to avoid user error), we set that to μ, the time-step of the $d\Gamma$ shock. For each simulated Δ_t, each worker i, and each job offer they have, k, we simulate the number $d\Gamma$ shocks they face in that product line as

$$P(d\Gamma = n) = \left(\frac{\Delta_t}{\mu}\right)^n \frac{e^{-\frac{\Delta_t}{\mu}}}{n!}$$

where we assume there can be no more than 20 jumps. This constructs a discrete CDF from which we can draw empirical jump numbers. We assume that the jumps across job offers are independent.

For each $d\Gamma$ they draw for each job over time step Δ_t, we simulate from the ODE given by C.77 to get the next state. These simulated draws are no longer on the discrete grid, so to get exact per-efficiency-unit wages we interpolate over the log wage / log quality distribution for both skill types using Matlab's *pchip* function. We also enforce the reflecting barrier, as we have in our transition matrix, in order to make things tie out.

Each worker has by K job offers at each time which have wages per skill unit $w(\hat{Q}_{ik})$. They also have a number of skill units per job offer given by equations 22 and 34 in the main text. In the discrete time simulations, we approximate the OU process with an AR(1). At any given point in time, a worker chooses job k out of K offers to maximize take-home pay. This gives us two arrays of dimension $N_w \times T \times K$, where $T = \frac{8}{\Delta_t}$ is the number of periods required to compute our growth rates.

This gives us the path of employment and wages in the absence of additional shocks. As discussed in the main text, we use 5 year forward-looking income compared to the three years of past income. The extra shock we consider is an additional $d\Gamma$ shock that happens instantly at the end of the 3rd year of employment.

We consider each possible outcome from this the worker's income separately and weight them appropriately. This is to say we consider that conditional on a $d\Gamma$ shock, there is a discrete jump distribution from 0 up to a K_{max}, and for each of these outcomes we compute the expected counterfactual path of wages for each worker. For resets, we do the same.

As in the data, we winsorize our outcomes at the 1st and the 99th percentile. To compute these expected quantiles of outcomes, we compute the CDF of potential outcomes given a $d\Gamma$ shock. For a product line outcome, o_l, conditional on a shock, s_l, let $P(s_l)$ be the probability that that type of shock occurs. Let $P(\omega_l)$ be the probability of state ω_l in the discretized product line stationary distribution.

Then the PDF of an aggregate outcome for product line l, o_l is given by

$$P(o_l) = P(s_l)P(\omega_l)$$

which is a discrete distribution because of the discrete steady state distribution and the finite set of outcomes conditional on the shocks. The CDF, $P(o_{f,p} \leq O)$ is given by sorting and integrating over the discrete distribution. The inverse CDF immediately gives the quantiles of the product line outcomes.

Workers vary not only by initial states, but also by product line level human capital z_{il}. Let $P(z_{il}|\omega_l)$ be the pdf of human capital for worker i in product line j be the density of the human capital

distribution given the stochastic process described in the main text for z_{ilt}. We approximate the distribution of $P(z_{il}|\omega_l)$ through standard simulation methods, drawing 15,000 simulated workers, and approximating the PDF with the empirical distribution of the simulations. Let $\hat{P}(z_{il}|\omega_l)$ be the realization of that empirical approximation, which gives PDF

$$P(o_{i,p}) \approx \hat{P}(z_{ij}|\omega_l)P(\omega_l)P(s_l).$$

The CDF follows immediately, as do the quantiles.

We take those winsorized product-line and worker-level outcomes, and use them to compute regression coefficients for both firms and workers. All variables are winsorized in the following sections. In the data, the outcomes are given by $dA_f = \frac{dV_f \mathbb{I}(dV_f>0)}{V_{f,t}}$ where $dV_{f,t} = V_{f,t} - V_{f,t-}$ because we cannot identify impacts at the product-line level. However in the model, we can observe product lines directly, and

$$dA_{f,l,t} = \frac{\Delta V_{f,l,t} \mathbb{I}(\Delta V_{f,l,t} > 0)}{V_{f,l,t}}$$

where $\Delta V_{f,l,t} = V_{f,l,t} - V_{f,l,t-}$. To compute $dA_{f,t}$ we sum over product lines,

$$dA_{f,t} = \sum_{l=1}^{P} dA_{f,l,t}$$

Competitor innovation conditional on a $d\Gamma$ shock is simple in this model. Since firms are infinitesimal it collapses to expected firm-level $dA_{f,t}$:

$$dA_i = \int_0^1 dA_{k,t} dk = E(dA_{f,t}|d\Gamma)$$

In the absence of a $d\Gamma$ shock the firm state does not change so both measures are equal to zero. This result is a key reason for using the machinery of the subordinated Γ process, because in continuous time if there were no innovation bursts, dAi would be a constant drift.

D.5 Computing Moments

For firms, we are interested in changes in five-year-forward expected growth paths of wage bill, profits, and employment. We have expressions that give instantaneous profit and employment given the current state, denote that vector y. To approximate a 5 year level of profits, we compute the forward-looking discounted value of y, but depreciate them at a 20% rate:

$$Y_{5yr} = (r - \Lambda)^{-1} y$$

where r is a diagonal matrix with 0.2 on the diagonal and Λ is the generator matrix for the stochastic process for the product line state. Instantaneous profits, employment, and pay follow from the numerical solution to the firm's problem described above.

For wage growth, we simulate forward a 5 year growth path for workers and compute log differences relative to a prior 3 years of employment. For the paths with an extra shock, described above, we do the same. The counterfactual wage growth path is the difference between these two outcomes.

The target moments of interest are largely regression coefficients of both firm and worker level outcomes on firm and competitor innovation. For some outcome, Y, we compute these directly as

$$\beta_o = E(X^T X)^{-1} E(X^T \tfrac{dY}{Y})$$

where all that varies is the outcome. We calculate $X^T X$ as

$$E(X^T X) = E \begin{pmatrix} dA_i^2 & dA_i dA_f \\ dA_i dA_f & dA_f^2 \end{pmatrix}$$

and

$$E(X^T Y_o) = E \begin{pmatrix} dA_i \frac{dY}{Y} \\ dA_f \frac{dY}{Y} \end{pmatrix}$$

Since $dA_i = E(dA_f)$ as described above is a constant, we can compute this once. However, it still remains to compute $E(dA_f)$ and $E(dA_f^2)$, and each $E(\frac{dY}{Y})$ and $E(dA_f \frac{dY}{Y})$.

Let $\hat{Q}_f = \{\hat{q}_{l,f}\}_{l=1}^{N_{lines}}$ be the set of product lines that the firm has. $dA_l = \frac{dV_l}{V_l}$, $dA_f = \frac{dV_f}{V_f} = \frac{V_{f,l}}{V_f} dA_l$, $A_f = \sum w_{f,l} a_{f,l}$; $w_l = \frac{V_{f,l}}{V_f}$ We need to compute:

$$E(dA_f | d\Gamma = \mu) = \frac{1}{N_{firms}} \sum_{f=1}^{N_{firms}} \sum_{l=1}^{N_f} w_{f,l} E(dA_{f,l} | \hat{q}_{f,l}, d\Gamma = \mu)$$

The expectation of the product line-level shock to the power k at the product line l for firm f given the quality $\hat{q}_{f,l}$, $E(dA_{f,l}^k)$, is given by

$$\begin{aligned} E(dA_{f,l}^k | \hat{q}_{f,l}) &= p(\text{Lose Line}) E(dA_{f,l}^k | \text{Lose Line}, \hat{q}_{f,l}) \\ &+ p(\Delta \hat{q}_{f,l} = -\mu \tilde{g}) E(dA_{f,l}^k | \Delta \hat{q}_{f,l} = -\mu \tilde{g}, \hat{q}_{f,l}) \end{aligned}$$

$$
\begin{aligned}
&+ \quad p(\text{Gain Line})E(dA_{f,l}^k|\text{Gain Line}, \hat{q}_{f,l}) \\
&+ \quad p(\text{Improve Line})\sum_{j=1}^{N} p(n=j|\text{Improve Line})E(dA_{f,l}^k|n, \text{Improve Line}, \hat{q}_{f,l})
\end{aligned}
$$

for $k = 1$.

For k=2 we also must include a cross term. Second moment computations become

$$E(dA_{f,l}^2|\hat{q}_{f,l}) = E(dA_{f,l,existing}^2|\hat{q}_{f,l}) + 2E(dA_{f,l,existing}|\hat{q}_{f,l})E(dA_{f,l,new}) + E(dA_{f,l,new}^2)$$

Let Y_f be another firm-level outcome, e.g. profit, employment, or pay. dY_f is the instantaneous change in that outcome in real time. Again, for simplicity, this is all conditioning on the shock, $d\Gamma = \mu$. The expectation terms which appear in the regression are

$$E\begin{pmatrix} dA_i \frac{dY_f}{Y_f} \\ dA_f \frac{dY_f}{Y_f} \end{pmatrix} = \begin{pmatrix} E(dA_i)E(\frac{dY_f}{Y_f}) \\ E(dA_f \frac{dY_f}{Y_f}) \end{pmatrix}$$

are what is left to compute. Thus we need $E(\frac{dY_f}{Y_f})$ and $E(dA_f \frac{dY_f}{Y_f})$.

$$E(\frac{dY_f}{Y_f}|d\Gamma = \mu) = \sum_{s \in S} p(\hat{q}_{f,l} = \hat{q}(s)) \frac{1}{N_{firms}} \sum_{f=1}^{N_{firms}} \sum_{l=1}^{N_f} w_{f,l}^* E(\frac{dY_f}{Y_f}|\hat{q}_{f,l}, d\Gamma = \mu)$$

where $w_{f,l}^* = \frac{Y_{f,l}}{Y_f}$ for a given outcome. Similar to $dA_{f,l}$,

$$
\begin{aligned}
E(\tfrac{dY_{f,l}}{Y_{f,t}}|\hat{q}_{f,l}) \quad &= \quad p(\text{Lose Line})E\left(\tfrac{dY_{f,l}}{Y_{f,t}}|\text{Lose Line}, \hat{q}_{f,l}\right) \\
&+ \quad p(\Delta\hat{q}_{f,l} = -\mu\tilde{g})E\left(\tfrac{dY_{f,l}}{Y_{f,t}}|\Delta\hat{q}_{f,l} = -\mu\tilde{g}, \hat{q}_{f,l}\right) \\
&+ \quad p(\text{Gain Line})E\left(\tfrac{dY_{f,l}}{Y_{f,t}}|\text{Gain Line}, \hat{q}_{f,l}\right) \\
&+ \quad p(\text{Improve Line})\sum_{j=1}^{N} p(n=j|\text{Improve Line})E\left(\tfrac{dY_{f,l}}{Y_{f,t}}|n, \text{Improve Line}, \hat{q}_{f,l}\right)
\end{aligned}
$$

$$E(dA_f dA_i) = E(dA_f)E(dA_i)$$

because firm f is atomistic.

$$E(dA_f^2|\hat{q}_{f,l}) = \frac{1}{N_{firms}} \sum_{f=1}^{N_{firms}} \left(\sum_{l=1}^{N_f} w_{f,l}^2 E(dA_{f,l}^2|\hat{q}_{f,l}) + \sum_{l=1}^{N_f} \sum_{m \neq l} w_{f,l} w_{f,m} E(dA_{f,l}|\hat{q}_{f,l})E(dA_{f,m}|\hat{q}_{f,m}) \right)$$

$$E\left(dA_f \frac{dY_f}{Y_f} | \hat{q}_{f,l}\right) = \frac{1}{N_{firms}} \sum_{f=1}^{N_{firms}} \Bigg(\sum_{l=1}^{N_f} w_{f,l} w^*_{f,l} E[dA_{f,l} \frac{dY_{f,l}}{Y_{f,l}} | \hat{q}_{f,l}] + \sum_{l=1}^{N_f} \sum_{m \neq l} w_{f,l} w^*_{f,m} E(dA_{f,l} | \hat{q}_{f,l}) E(\frac{dY_{f,m}}{Y_{f,m}} | \hat{q}_{f,m}) \Bigg)$$

To compute the unconditional expectations, we just integrate over the stationary distribution for the probability that product line l has quality $\hat{q}$.

$$E\left(dA_f \frac{dY_f}{Y_f}\right) = \sum_{s \in S} p(\hat{q}_{f,l} = \hat{q}(s)) E\left(dA_f \frac{dY_f}{Y_f} | \hat{q}_{f,l}\right),$$

which we can do analogously for any of these outcomes, following from the definition of conditional expectation.

For wage growth, we have an additional complication, which is that we are integrating over a different distribution of $\hat{q}_{f,l}$ when we condition on income. Rather than using the steady state product line distribution, we use the observed $\hat{q}_{f,i \in l}$ (where i indexes the worker) for a single product line in the firm, and the remainder are simulated from the product line stationary distribution. The firm's other product lines are drawn from the stationary distribution.

For workers, their wage growth only depends on the $\hat{q}_{i \in l}$, so

$$E(\frac{dW_i}{W_i} dA_{f,i \notin l} | \hat{q}_i, \hat{q}_l) = E(\frac{dW_i}{W_i}) E(dA_{f,l}).$$

While before we had a vector of product line outcomes, $\frac{dY_f}{Y_f}$ and a vector of product line dA_f shocks, in this case we have a scalar and a vector. Thus our above formula simplifies conditional on $\hat{q}_{f,i \in l}$,

$$E(\frac{dW_i}{W_i} dA_f | \hat{q}_{f,i \in l}) = \frac{1}{N_{firms}} \sum_{f=1}^{N_{firms}} \left[w_{f,l} E\left(\frac{dW_i}{W_i} dA_{f,i \in l} | \hat{q}_{f,l}\right) + E\left(\frac{dW_i}{W_i}\right) \sum_{m \neq l} w_{f,m} E(dA_{f,m} | \hat{q}_{f,m}) \right].$$

This integrates over the possible firms that a product line containing a worker, i working in product line l with quality $\hat{q}_{f,l}$ could have. We then need to integrate over the worker's $\hat{q}$ distribution, which we do by averaging across simulated workers,

$$E\left(\frac{dW_i}{W_i} dA_f\right) \approx \frac{1}{N_{workers}} \sum_{i=1}^{N_{workers}} E(\frac{dW_i}{W_i} dA_f | \hat{q}_{f,i \in l})$$

To condition on income, we just subset the worker distribution such that a worker meets a certain criterion for starting income.

The computation associated with looping over the quadratic set of cases, $l = m$ and $l \neq m$ is extremely costly, especially for firms a large number of product lines, which causes it to explode combinatorically. However, we can simplify the computation by boiling it down to a portfolio of two product lines: one which we take from the focal state we are conditioning on, and one which we take as an existing portfolio of product lines in the rest of the firm. Our conditioning information is only in a single one of the firms product lines, denoted $\hat{q}_{f,l}$, which we know is equal to $\hat{q}(s)$.

First, note that

$$E(dA_f^2|\hat{q}_{f,l}) = E\left(w_{f,l}^2 dA_{f,l}^2|\hat{q}_{f,l}\right) + \sum_{l\neq m}\left[E(w_{f,l}w_{f,m}dA_{f,l}dA_{f,m}|\hat{q}_{f,l}) + E(w_{f,m}^2 dA_{f,m}^2|\hat{q}_{f,l})\right]$$

Importantly, $dA_{f,m}$ and $dA_{f,l}$ are independent in this construction, so $dA_{f,m}^2$ does not depend on the value of $\hat{q}_{f,l}$ and the conditional expectation is equal to the unconditional. Consider the terms not involving the focal product line, l. They condition on $\hat{q}_{f,l}$ *only* through the denominator of $w_{f,l}$, V_f. Let $V_{f,-l} = \sum_{m\neq l} V_{f,m}$ be the leave-one-product-line out value of the firm. Identically, the change in firm value associated with an event in product line l is

$$dA_f = w_{f,l}dA_{f,l} = \frac{V_{f,l}}{V_f}\frac{dV_{f,l}\mathbb{I}(dV_{f,l} > 0)}{V_{f,l}}$$

For all other product lines, then, we can write

$$dA_f = w_{f,m}dA_{f,m} = \frac{V_{f,m}}{V_f}\frac{dV_{f,m}\mathbb{I}(dV_{f,m} > 0)}{V_{f,m}} = \frac{V_{f,-l}}{V_f}\frac{V_{f,m}}{V_{f,-l}}\frac{dV_{f,m}\mathbb{I}(dV_{f,m} > 0)}{V_{f,m}}$$

This means we can collapse all other product lines into a single $dA_{f,-l}$ and only consider a synthetic "two-product-line" firm for our simulations, speeding up computation considerably relative to looping over each product line and computing the product. All we need to do is track the total share of firm value associated with the non-focal product line, and use that as a normalizing factor. In other words, for each firm f we only need

$$dA_{f,-l} = \frac{\sum_{m\neq l} dV_{f,m}\mathbb{I}(dV_{f,l} > 0)}{V_{f,-l}}$$

and $V_{f,-l}$, which we can compute directly from our firm simulations. Second moment computations become

$$E(dA_f^2|\hat{q}_{f,l}) = w_{f,l}^2 E(dA_{f,l}^2|\hat{q}_{f,l}) + 2w_{f,l}w_{f,-l}E(dA_{f,l}|\hat{q}_{f,l})E(dA_{f,-l}) + w_{f,-l}^2 E(dA_{f,-l})^2$$

and for first moments we have

$$E(dA_f|\hat{q}_{f,l}) = w_f E(dA_{f,l}|\hat{q}_{f,l}) + w_{f,-l}E(dA_{f,-l})$$

for other outcomes it becomes

$$\begin{aligned} E\left(dA_f \tfrac{dY_f}{Y}|\hat{q}_{f,l}\right) = w^y_{f,l}w^v_{f,l}E\left(dA_{f,l}\tfrac{dY_{f,l}}{Y}|\hat{q}_{f,l}\right) + w^v_{f,l}w^y_{f,-l}E(dA_{f,l}|\hat{q}_{f,l})E\left(\tfrac{dY_{f,-l}}{Y_{f,-l}}\right) \\ + w^v_{f,-l}w^y_{f,l}E(dA_{f,-l})E\left(\tfrac{dY_{f,l}}{Y_{f,l}}|\hat{q}_{f,l}\right) + w^v_{f,-l}w^y_{f,-l}E\left(dA_{f,-l}\tfrac{dY_{f,-l}}{Y_{f,-l}}\right) \end{aligned}$$

All that we need to approximate, then, is the distribution of firm weights $w_{f,l}$ and $w_{f,-l}$ given the value of the focal product line, l. We do this by simulation, drawing 3,000 firms which we use to compute the distribution of weights. The expectations integrated over firms are computed numerically, such that (for example, in the case of a firm-level covariance with dA_f)

$$E\left(dA_f \tfrac{dY_f}{Y}|\hat{q}_{f,l}\right) \approx \frac{1}{N_{sims}} \sum_{\hat{f}=1}^{N_{sims}} E\left(dA_{\hat{f}} \tfrac{dY_{\hat{f}}}{Y}|\hat{q}_{\hat{f},l}\right)$$

where each firm $\hat{f}$ is drawn from the distribution of firms.

To simulate firms, we draw from the firm size distribution described in equation C.63. For each product line we draw, we then simulate a quality from the discrete approximation to the stationary distribution. In order to match our Compustat sample, we mirror the censoring that happens in the data. A competitor shock type for a single product line firm eliminates that firm from the forward-looking sample. We account for this in the model by omitting that possibility from the sample, and re-weighting the prevalence of single- and multiple-product line firms to account for the censoring.

Wages have one additional step, which comes from the fact that we have two skill types. Let $\frac{dW_i}{W_i}$ be the percent wage growth of worker i. That worker has a vector of human capital draws, z_i, of size K. In computing regression coefficients, we need to (optionally conditional on lagged income group) integrate over the distribution of human capital and available skill types.

The path of human capital is drawn from the same random normal process, but plugged into the inverse CDF of a Frechet distribution with parameters T_l and θ_l for skill type l and parameters T_m and θ_m for skill type m. This preserves the respective skill type distributions while saving simulation time. Workers are grouped into income bins based on the quantiles of their average three year income in the CDF of workers.

We need to be able to integrate over two distinct skill types. To do this, we compute the empirical

share of the number of simulated workers, weighted by the share of the type in the economy, in each income bin. For each skill type, we compute

$$E\left(\frac{dW_i}{W_i}dA_f\right) \approx \frac{1}{N_{workers}} \sum_{i=1}^{N_{workers}} E(\frac{dW_i}{W_i}dA_f|\hat{q}_{f,l}(i))$$

where $\hat{q}_{f,l}(i)$ is the state associated with the worker at the time of the shock. We compute this separately for each type, and then average them based on the empirical shares. Let ω_l be the share of workers with skill type l and $\omega_m = 1 - \omega_l$ be the share associated with m. Then for a given income group, g, we first compute the total measure of workers. Let $s(i)$ be the skill type of worker i. Then the total measure of workers in an income group is $\sum_{i \in g} \omega_{s(i)}$ and the relative share of type l is given by $w_g(l) = \frac{\sum_{i \in g, s(i)=l} \omega_{s(i)}}{\sum_{i \in g} \omega_{s(i)}}$.

$$\beta_g = \left(w_g(l)E(X_g(l)^T X_g(l)) + w_g(m)E(X_g(m)^T X_g(m))\right)^{-1} \left(w_g(l)E(X_g(l)^T Y_g(l)) + w_g(m)E(X_g(m)^T Y_g(m))\right)$$

where the expectations are approximated by the empirical distribution of simulated workers, and calculated as described above.

The standard deviation of wage growth is equal to

$$\sigma(wg) = \sqrt{\omega_m \hat{\sigma}_m^2 + \omega_l \hat{\sigma}_l^2}$$

where $\hat{\sigma}^2$ is the empirical variance of wage growth rates for the simulated worker panel for each type.

For each income group we compute the share of total pay that goes to the group. For a group g the total income paid to that group is given by

$$I(g) = \omega_l \sum_{i \in g, s(i) \in l} z_{ij} w(\hat{q}_j) + \omega_m \sum_{i \in g, s(i) \in m} z_{ij} s(\hat{q}_j)$$

The share paid to each group is simply

$$\frac{I(g)}{\sum_{g' \in G} I(g')}.$$

Three year exit probabilities are computed based on the empirical frequency with which individuals switch jobs in the simulated panel, appropriately weighted by type.

D.6 Computing the objective function

The loss function is the traditional method of simulated moments quadratic form,

$$\mathcal{L}(\hat{\theta}) = g(\hat{\theta})' W g(\hat{\theta})$$

W is diagonal, and is fixed through the estimation in order to emphasize the fit of certain moments are more important than others. Because the moments are of different scales, we first scale them by their empirical counterparts before computing the loss.

We optimize the objective function using Matlab's surrogateopt, which performs radial basis function expansion for global optimization.

D.7 Computing Certainty Equivalent Welfare

Our measure of certainty equivalent welfare is the willingness to pay, annually for 20 years, for a subsidy that increases innovation by a rate that leads to a 1 percentage point increase in the growth rate for a single year. Our baseline growth rate, g is set to 2%, our time discount factor, β, is set to 4% in order to be two percentage points above the growth rate, and the subsidy increases the growth rate, temporarily, by a factor Δ, which we set to 1%. For convenience of notation we suppress worker indexes here, as everything is at the level of the individual.

Worker preferences are such that given a wage rate they maximize effort:

$$\max_{e} u\left(e_{k,i,t} z_{j(k,i,t)} Y(t) w_{k,t} - Y(t)\frac{e^{1+\varphi}}{1+\varphi}\right)$$

so at optimal effort e^*

$$u(w, z, e^*) = Y_t^{1-\gamma} u\left(\frac{\varphi}{1+\varphi}[wz]^{1+\frac{1}{\varphi}}\right)$$

We make the assumption that workers are hand-to-mouth, and thus eat their wages at each time t, so consumption c_t is

$$c_t = \frac{\varphi}{1+\varphi}[wz]^{1+\frac{1}{\varphi}}$$

We work with CRRA preferences with risk aversion $\gamma \in \{1, 2\}$, so

$$u(c_t) = \frac{c_t^{1-\gamma}}{1-\gamma}$$

Note, here, that aggregate growth does not enter in a risk neutral fashion. We want to find a scalar, proportional consumption adjustment factor, a such that people are indifferent between the alternative growth path and a wage path multiplied by this constant factor.

$$E_0\left[\int_0^{20} e^{-(\beta-g(1-\gamma))t}u(a\times c_t)dt\right] = E_0\left[\int_0^{1} e^{-(\beta-(g+\Delta)[1-\gamma])t}u(\tilde{c}_t)dt + \int_1^{20} e^{-(\beta-g[1-\gamma])t+\Delta(1-\gamma)}u(\tilde{c}_t)dt\right] \tag{C.79}$$

where $\tilde{c}_t$ is the counterfactual consumption for the worker under the subsidy.

We can factor $a^{1-\gamma}$ out of the integrals on the left hand side front since the utility function is homogenous of degree $1-\gamma$.

$$\int_0^T e^{-(\beta-g(1-\gamma))t}u(a\times c_t)dt = a^{1-\gamma}\int_0^T e^{-(\beta-g(1-\gamma))t}u(c_t)dt$$

So we can factor $a^{(1-\gamma)}$ and get

$$a_i = \left(\frac{E_0\left[\int_0^1 e^{-(\beta-(g+\Delta)[1-\gamma])t}u(\tilde{c}_t)dt + \int_1^{20} e^{-(\beta-g[1-\gamma])t+\Delta(1-\gamma)}u(\tilde{c}_t)dt\right]}{E_0\left[\int_0^{20} e^{-(\beta-g(1-\gamma))t}u(c_t)dt\right]}\right)^{\frac{1}{1-\gamma}}$$

where we evaluate by simulation. We consider three counterfactuals in the text, all of which raise growth by the same amount. First, we raise the aggregate innovation rate proportionally, so that

$$\frac{I_{cf}}{I_{cf}+x_{cf}+e_{cf}} = \frac{I}{I+\bar{x}+\bar{e}}$$

and

$$I_{cf}+x_{cf}+e_{cf} = b[I+\bar{x}+\bar{e}]$$

where a factor b is chosen so that $\Delta = 0.01$. In the second and third counterfactuals we keep the second identity, but attribute it entirely to I and entirely to $\bar{x}+\bar{e}$ instead of keeping the rates proportional to the calibration.

References

Acemoglu, D., U. Akcigit, H. Alp, N. Bloom, and W. Kerr (2018, November). Innovation, reallocation, and growth. *American Economic Review 108*(11), 3450–91.

Acemoglu, D., G. Gancia, and F. Zilibotti (2012). Competing engines of growth: Innovation and standardization. *Journal of Economic Theory 147*(2), 570–601.e3.

Chow, M. C., T. C. Fort, C. Goetz, N. Goldschlag, J. Lawrence, E. R. Perlman, M. Stinson, and T. K. White (2021). Redesigning the longitudinal business database. Technical report, National Bureau of Economic Research.

De Loecker, J., J. Eeckhout, and G. Unger (2020, 01). The Rise of Market Power and the Macroeconomic Implications*. *The Quarterly Journal of Economics 135*(2), 561–644.

Gouin-Bonenfant, E. (2022). Productivity dispersion, between-firm competition, and the labor share. *Econometrica 90*(6), 2755–2793.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of us workers reveal about lifecycle earnings dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., S. Ozkan, and J. Song (2014). The Nature of Countercyclical Income Risk. *Journal of Political Economy 122*(3), 621–660.

Hoberg, G. and G. Phillips (2016). Text-based network industries and endogenous product differentiation. *Journal of Political Economy 124*(5), 1423–1465.

Jarmin, R. S. and J. Miranda (2002). The longitudinal business database. *Available at SSRN 2128793*.

Klette, T. J. and S. Kortum (2004). Innovating firms and aggregate innovation. *Journal of Political Economy 112*(5), 986–1018.

Kogan, L., D. Papanikolaou, A. Seru, and N. Stoffman (2017). Technological Innovation, Resource Allocation, and Growth. *The Quarterly Journal of Economics 132*(2), 665–712.

Nadarajah, S. (2009). Some truncated distributions. *Acta applicandae mathematicae 106*(1), 105–123.

Peters, M. (2020). Heterogeneous markups, growth, and endogenous misallocation. *Econometrica 88*(5), 2037–2073.

Schmidt, L. D. and Y. Zhu (2016). Quantile spacings: A simple method for the joint estimation of multiple quantiles without crossing.

Seegmiller, B. (2024). Valuing labor market power: the role of productivity advantages. Working paper, Northwestern University.