

Market Concentration, Growth, and Acquisitions*

Joshua Weiss[†]

IIES

November 11, 2023

[Click here](#) for the most recent version

Abstract

I develop an oligopolistic growth model in which a firm chooses how much to innovate, as well as the degree to which its innovation is directed toward its competitors' goods through creative destruction. I find that a firm's size shapes the direction of its innovation: larger firms generate less growth relative to the rate at which they creatively destroy their competitors' goods. I demonstrate positive and normative implications of this mechanism. First, an increase in large firm innovation incentives can explain a substantial portion of the recent growth slowdown in the US. Second, it is optimal to tax large firm revenues because the direction of innovation is more significant than other size-related distortions. Third, a tax on large firm acquisitions of smaller competitors' goods can reduce welfare by encouraging large firms to innovate rather than acquire. Fourth, it is advantageous to be large.

Keywords: Market Power, Industry Concentration, Innovation, Creative Destruction, Arrow Replacement Effect, Acquisitions

*I thank Timo Boppart, Maarten De Ridder, Basile Grassi, Per Krusell, Kurt Mitman, Paula Onuchic, Pau Roldan-Blanco, and Dajana Xhani for helpful comments and feedback.

[†]Email: joshua.weiss@iies.su.se

1 Introduction

Within industries at the national level in the US, the largest firms have steadily increased their shares of economic activity—in particular, revenue and innovation—since the 1990s.¹ Over this period, the productivity growth rate rose dramatically, but then fell below where it started.² It is thus important to understand how a firm’s size within its industry shapes its economic decisions, and the macroeconomic consequences. This can help us make sense of recent changes in the US economy, and inform optimal size-dependent and antitrust policies.

In this paper, I develop a model to study the implications of the joint concentration of sales and innovation. The novel feature of the analysis is that a firm chooses not only the magnitude, but also the direction of its innovation: the degree to which its innovations are close substitutes with its competitors’ goods. The direction of a firm’s innovation determines how much it generates growth, which benefits consumers through lower prices, and how much it harms industry competitors by depreciating the value of their past innovations. Depreciation occurs through two channels. First, industry-wide growth erodes a good’s sales over time by pushing up its relative price. For example, if soy milk improves, then sales of oat milk and almond milk fall. Second, a good can be improved on by a close substitute—creatively destroyed—in which case it quickly loses a substantial share of its sales. If a soy milk producer develops a version of oat milk, then to some extent oat milk as a whole improves and so almond milk loses sales due to growth, but other oat milk producers face the biggest losses. These two forms of depreciation have a significant impact on a firm’s value of an innovation: using US Census data, Garcia-Macia, Hsieh, and Klenow (2019) estimate that from 1983-1993, erosion from growth and replacement by creative destruction depreciated a good’s employment by 5% and 7% on average per year, respectively.³

The key finding is that taking into account its direction, large firm innovation is especially harmful to industry competitors. Specifically, the larger a firm’s share of industry revenue, the lower the rate at which it generates growth relative to the rate at which it creatively destroys each of its competitors’ goods. Thus, if sales and innovation are concentrated within

¹See Olmstead-Rumsey (2022) for the concentration of revenue and R&D expenditures in Compustat data. See Weiss (2020) for the concentration of revenue in US Census data.

²See Garcia-Macia, Hsieh, and Klenow (2019) for the growth rate, and Aghion, Bergeaud, Boppart, Klenow, and Li (2022) and De Ridder (2021) for the joint dynamics of concentration and the growth rate.

³Using retail scanner data from 2006-2015, Argente, Lee, and Moreira (2021) estimate that a good’s sales peak after a year, and then depreciate by 30% on average per year.

few firms, then their competitors face a high depreciation rate relative to the growth rate.

This mechanism links the recent rise in industry concentration to the productivity growth boom and then slowdown. A positive shock to large firms, for example to their profitability, pushes them to innovate faster. Concentration and growth increase. But their competitors respond by innovating less, which lowers growth. Growth can fall below where it started because the rise in large firm innovation implies a particularly big increase in the rate at which their competitors' goods are creatively destroyed, and thus a particularly big response by those competitors. More generally, inserting oligopoly and the direction of innovation into other theories of the rise in concentration and growth slowdown will yield a bigger slowdown.

A corollary is that taxing large firm activity can increase growth and welfare relative to a tax-free benchmark. Intuitively, this has the opposite effect of a rise in large firm profitability.

Finally, I highlight two implications of the theory for antitrust policy, which are particularly relevant given a recent dramatic rise in the rate at which venture capital backed startups are acquired by public firms relative to the rate at which they go public.⁴ First, large firm acquisitions of smaller competitors' goods present a novel trade-off for policy makers. On the one hand, these acquisitions increase the concentration of revenue, which further distorts the direction of large firm innovation and thus decreases growth. On the other hand, the expectation of high value acquisitions encourages large firms to acquire rather than innovate, which can lead to faster growth. Second, large firms are particularly inclined to acquire their competitors' innovation technologies and goods, i.e., there is a high return to size. Large firms protect their own goods while deterring competition by choosing a higher creative destruction rate for their competitors' goods than for their own.

Formally, I introduce a growth model with two key components. First, it is oligopolistic. There is ex-ante permanent heterogeneity across firms in innovation costs; for simplicity, each industry consists of a single "large" firm with a sufficiently low innovation cost so that it has a discrete effect on industry aggregates, and a mass of "small" firms each of which has a negligible effect on industry aggregates. A large firm conducts a non-negligible share of its industry's innovation, so how it innovates is relevant. As a result, a large firm earns a non-negligible share of its industry's revenue, which shapes the direction of its innovation.

The second key model component is that firms innovate by improving on and replacing their own or their competitors' goods (creative destruction) and by internally innovating,

⁴Documented in Pellegrino (2021).

which can mean improving their own goods or developing new ones. All innovation generates growth, which benefits consumers and takes sales from old goods. Creatively destroying a competitor's good entails additional business stealing unrelated to growth because intuitively, even if the improvement is negligible, nearly indifferent consumers switch from the old version to the new one, yielding a big shift in sales with little effect on real consumption.

The important outcome of these model features is a negative relationship between a firm's industry revenue share and the rate at which the firm generates growth relative to the rate at which it creatively destroys each of its industry competitors' goods. This follows from two mechanisms. First, a firm's industry revenue share shapes the relative rates at which it engages in different types of innovation: a higher share weakens its incentive to internally innovate or creatively destroy its own goods relative to its incentive to creatively destroy its competitors' goods. This relates to the Arrow (1962) replacement effect that incumbents' incentives to innovate are lower due to fears of cannibalizing the sales of their past innovations; or to the force in static models of oligopolistic competition such as Edmond, Midrigan and Xu (2022) that pushes larger firms to set higher markups. Here, cannibalization affects *relative* innovation rates because generating growth takes sales equally from all old goods, many of which a large firm produces, whereas the additional business stealing from creatively destroying a competitor's good takes sales only from that competitor.

The second mechanism concerns only creative destruction. Firms prefer to creatively destroy their competitors' goods rather than their own. Thus, a larger firm generates less growth through creative destruction relative to the rate at which it creatively destroys each of its competitors' goods because those goods comprise a smaller share of the industry.

These mechanisms follow only from a firm minimizing the innovation cost of a given rate at which it gains revenue share. Neither entails a large firm choosing the direction of its innovation *in order to* deter competition. A large firm does not have commitment, and its competitors' decisions depend on their expectation of its future innovation.

The main experiment I study in the model is a permanent rise in large firm innovation incentives. This captures many shocks or policy changes that only directly affect large firms, such as a rise in their profitability or span of control, a fall in their innovation costs, or a fall in the tax rate on their revenues. I find a sufficient statistic in terms of observable equilibrium outcomes for whether long-run growth falls. In particular, growth falls if industries are sufficiently concentrated, if the small firm innovation rate is sufficiently elastic with respect

to the value of innovating, and if creative destruction risk is sufficiently important for that value. Moreover, I show that if variation in large firm innovation incentives drives variation in outcomes across industries, then growth as a function of concentration is an inverted-U.

To evaluate the quantitative significance of the mechanisms, I calibrate the model. In particular, I use industry concentration data and large negative firm-level employment flows from Garcia-Macia, Hsieh, and Klenow (2019) to calibrate the size of large firms and the creative destruction risk firms face, respectively. I find that a permanent decrease in large firm per-good fixed costs—calibrated to match the rise in concentration since the 1990s—explains 35% of the long-run growth slowdown. Growth due separately to creative destruction and internal innovation fall, as estimated in the data. The experiment also generates significant shares of the observed short-run burst in growth, decline in R&D efficiency, positive correlation between rising concentration and growth across industries, and decline in entry.

Taxing large firm revenues to undo the rise in concentration improves welfare. This stands in contrast to the idea that large firms under-produce or under-innovate, and that subsidizing them improves welfare by reducing misallocation. To demonstrate the greater quantitative significance of the direction of large firm innovation over misallocation concerns, I show that a tax is beneficial even if the rise in large firm innovation requires no R&D costs.

Next, I allow large firms to search for and make take-it-or-leave-it offers to acquire their competitors' goods, subject to a tax on the payments. I find a sufficient statistic in terms of observable equilibrium outcomes for whether a permanent increase in the tax leads to lower long-run growth. In particular, growth falls if the acquisition rate is sufficiently inelastic with respect to the value of each acquisition, if firms' innovation rates are sufficiently elastic with respect to the value of innovating, and if concentration is sufficiently high. Intuitively, large firms internalize that when they innovate, they depreciate the value of their competitors' goods, and so of future acquisition opportunities. Thus, a higher tax on acquisitions reduces the acquisition rate, but leads to more large firm innovation, which is particularly bad for growth if concentration is high.

Finally, I allow a large firm to acquire a small competitor's ability to internally innovate or creatively destroy. In the former case, there is no directional choice, and the large firm's value of the technology is less than the small firm's. Intuitively, the large firm lowers the innovation rate to limit cannibalization, and anticipating this, its competitors innovate more. In the latter case, there is a directional choice, and the large firm's value of the technology is

greater than the small firm's. It lowers the creative destruction rate for its own goods, but increases the rate for its competitors' goods, so anticipating this, its competitors innovate less. Thus, absent a directional choice, the return to size is negative, but with a directional choice, it is positive.

Empirical Studies of Large Firm Innovation: Cavenaile, Celik, and Tian (2021) estimate that across industries, a variety of innovation measures are inverted-U functions of industry concentration. This provides indirect evidence for the theory, which also generates an inverted-U. Moreover, whereas previous theories of the inverted-U are based on dueling effects of competition on a firm's incentive to innovate, the theory here is based on dueling effects of a large firm's success.⁵ A highly profitable large firm innovates quickly, but in a way that is particularly discouraging to its competitors, which leads to less overall innovation.

Akcigit and Kerr (2018) estimate that the larger a firm, the more their patents tend to cite their own past patents rather than their competitors'. This appears at odds with the mechanism in my model that the larger a firm, the lower the rate at which they generate growth relative to the rate at which they creatively destroy their competitors' goods. However, they rationalize their findings with a model in which a firm's ability to internally innovate scales with its size more than its ability to externally innovate. On the other hand, in my model, a firm's ability to innovate is fixed. As I discuss in Section 8, to align with their findings, I can extend the model so that a firm's ability to internally innovate scales with its share of past innovations. This does not weaken the results because what matters is how a firm innovates relative to its ability. Indeed, the model in Akcigit and Kerr (2018) underestimates the empirical internal innovation share for medium sized firms, but overestimates it for the largest firms. This suggests that relative to their ability, the largest firms direct their innovation toward their competitors, as my model predicts.

Another prediction of my model is that large firms shield their own goods from creative destruction, but continue to target their competitors'. Akcigit and Kerr (2018) and Akcigit, Alp, and Peters (2021) find empirical support. In the former, larger firms' patents receive fewer external citations. In the latter, a lower creative destruction rate for larger firms' goods helps explain the joint distribution of firm size and age.

Finally, Argente, Baslandze, Hanley, and Moreira (2021) document that larger firms

⁵Aghion, Bloom, Blundell, Griffith, and Howitt (2005) also estimate and explain the inverted-U.

are more likely to use patents to deter competition but not to develop products. This demonstrates that a firm's size shapes its innovation decisions. Moreover, our theories are complementary because creative destruction is most relevant where patent protection is weak.

The Rise in Concentration and Fall in Growth: In terms of explaining recent changes in the US economy, the most similar papers are Aghion, Bergeaud, Boppart, Klenow, and Li (2022) and De Ridder (2021). They also explain both the speedup and slowdown in growth that coincided with the rise in concentration. In their theories, all firms are infinitesimal, and a rise in innovation by high process productivity firms leads to slower growth. These firms are difficult to compete with, so their presence reduces the return to creative destruction. As such, I find similar results from the rise in large, but equally productive firms that they find from the rise in high productivity firms. The key difference is that my model generates a fall in growth due to each of internal innovation and creative destruction, whereas their models only incorporate creative destruction. Moreover, if they did incorporate internal innovation, then it would likely rise in their experiments because firms face lower rates of creative destruction. This difference is important because Garcia-Macia, Hsieh, and Klenow (2019) estimate that most of the fall in growth was due to internal innovation.

Size-Related Distortions: I contribute to a substantial literature on how a firm's size within its industry affects its economic decisions. Specifically, to the theory of large firms in settings without collusion or commitment. Shapiro (2012) discusses the literature concerning innovation, and Edmond, Midrigan, and Xu (2022) and Berger, Herkenhoff, and Mongey (2022) are recent examples in product and labor market settings, respectively. The point of departure is that these papers study environments in which decisions are one-dimensional, whereas in my model, a firm also chooses how much to direct its innovation toward creatively destroying its competitors' goods. In one-dimensional models, it is optimal to subsidize large firm activity because the key size-related distortion is that larger firms do too little. By contrast, I find that subsidizing large firm activity lowers welfare because the size-dependency of the direction of innovation is a more significant force.

It is useful to compare my results to those of Weiss (2020) in which large firms use intangibles to commit to higher production and deter competition. In that setting, large firms still under-produce, so their desire to deter competition improves welfare because they can

only do so by producing more. By contrast, in this paper, large firms do not choose the direction of their innovation in order to deter competition. Nonetheless, their choice lowers welfare because it increases the costs imposed on their competitors for a given innovation rate.

Innovation and Acquisitions: “Entry for buyout”, described in Rasmusen (1988) and recently in Letina, Schmutzler, and Seibel (2021) and Fons-Rosen, Roldan-Blanco, and Schmitz (2022), is another theory in which the expectation of future acquisitions increases growth: entrants innovate to be acquired. Yet, if large firm innovation is one-dimensional, then this effect needs to be weighed against the negative effect on large firm innovation. Indeed, Fons-Rosen, Roldan-Blanco, and Schmitz (2022) find that reducing acquisitions improves growth by increasing large firm innovation. In my model, acquisitions are made with take-it-or-leave-it offers, so entry for buyout is absent. Instead, acquisitions may increase growth precisely because they reduce large firm innovation. An important distinction is that if large firms innovate too little, then acquisitions are worse in more concentrated industries, whereas in my theory, acquisitions are better in more concentrated industries.

Model Building Blocks: I build on creative destruction models such as Grossman and Helpman (1991b), Aghion and Howitt (1992), and Klette and Kortum (2004), and expanding varieties models such as Romer (1990) and Grossman and Helpman (1991a). Atkeson and Burstein (2019) also combine the two, but without large firms. Finally, firms can improve on their own goods as in Peters (2020), but they do so to increase sales rather than markups.

The paper proceeds as follows. In Section 2, I develop the industry model. In Section 3, I discuss the main mechanisms and prove qualitative results. In Section 4, I embed the industry model in a macroeconomic model and calibrate. In Section 5, I conduct the main quantitative experiment. In Section 6, I analyze policy in the calibrated model. In Section 7, I study the value of size. In Section 8, I discuss key model assumptions. Section 9 concludes.

2 Industry Model

I introduce an oligopolistic industry model, made up of a static block which describes the within period production game, and a dynamic block consisting of firms’ decisions to innovate

and acquire other firms' goods. I analyze this industry model in Section 3. In Section 4.1, I embed it into a macroeconomic model, which I then use for the remainder of the paper. Although I focus on one industry in this section, I include a subscript n , which will later index industries, for industry specific variables so that it is clear which variables are determined at the macroeconomic level, i.e., are exogenous to the industry.

Time is continuous and indexed by $t \in [0, \infty)$. At each time t , the static block is as follows. There is a measure N_t of small firms and a single large firm. Small firms are indexed by $i \in [0, N_t]$, and a generic small firm is denoted by S ; the large firm is denoted by $i = L$. Firms produce a continuum of intermediate goods, which they sell to a representative intermediate good producer. Firms take as given the measure $M_{n,t}$ of goods, indexed by $j \in [0, M_{n,t}]$, and the productivities with which they can produce them. They hire labor and purchase a fixed cost input in perfectly competitive markets to produce and maximize static profits. The outcome is that for each good, there is a unique firm with positive sales for that good. Moreover, there is a single markup that all firms set for all goods. I discuss the importance of a single markup in Section 8.

In the dynamic block, firms hire labor to innovate and search for acquisition opportunities to maximize expected discounted profits. Firms innovate by creatively destroying old goods and internally innovating. All types of innovation generate growth by increasing the productivity of an old good or by expanding the set of goods. Creatively destroying a good currently produced by a competitor also shifts the unique producer of that good to the innovating firm. Only the large firm can acquire its competitors' goods, and doing so shifts the unique producer without affecting productivity.

The key difference between small firms and the large firm is the latter's innovation cost is infinitely lower, so it innovates on the same scale as the entire mass of small firms.

2.1 Static Block

2.1.1 Demand

At each time t , the industry good producer takes as given a price offered by each firm i for each good j , $\{p_{n,t}(i, j)\}$, and chooses demand to maximize an aggregate $Y_{n,t}$ subject to

spending R_t . The aggregate is linear within each good and CES across goods:

$$Y_{n,t}^{\frac{\gamma-1}{\gamma}} = \int_0^{M_{n,t}} y_{n,t}(j)^{\frac{\gamma-1}{\gamma}} dj,$$

where $y_{n,t}(j)$ is the sum of its demand for good j across all firms, and $\gamma > 1$ is the constant elasticity of substitution across goods. Since quantities from different firms are perfect substitutes, the industry good producer purchases only from firms that offer the lowest price, $p_{n,t}(j) \equiv \min_i \{p_{n,t}(i, j)\}$. Thus, the budget constraint is $\int_0^{M_{n,t}} p_{n,t}(j) y_{n,t}(j) dj = R_t$. The solution to this problem yields the demand curve⁶

$$y_{n,t}(j) = p_{n,t}(j)^{-\gamma} P_{n,t}^{\gamma-1} R_t,$$

where $P_{n,t}$ is the industry price index so that $R_t = P_{n,t} Y_{n,t}$, given by $P_{n,t}^{1-\gamma} = \int_0^{M_{n,t}} p_{n,t}(j)^{1-\gamma} dj$.

2.1.2 Production and Competition

At each time t , each firm can potentially produce each good using two different production technologies, each of which is linear in labor. First, firm i has an idiosyncratic version of each good j , which has productivity $z_{n,t}(i, j)$. Second, firm i can produce an imitation of the most productive version, with productivity $z_{n,t}(j)/\sigma$, where $z_{n,t}(j) \equiv \max_i \{z_{n,t}(i, j)\}$ is the maximum productivity at time t across all firms' idiosyncratic versions of good j , and $\sigma > 1$ is a discount that captures the ability of firms to imitate each others' technologies.

Production and competition occur in two stages at each time t . In the first stage, each firm chooses for which goods it will pay a fixed cost. If firm i pays the fixed cost for good j , then it can produce its idiosyncratic version. Otherwise, it can produce an imitation of the most productive version, regardless of whether that version's firm pays the fixed cost. In the second stage, fixed cost payments are common knowledge, and each firm chooses a price for each good. Firms then produce to satisfy industry good producer demand.

It will be convenient for fixed costs to scale in a particular way with a good's productivity. To that end, define $Z_{n,t}$ to be an industry aggregate of productivity:

$$Z_{n,t}^{\gamma-1} = \int_0^{M_{n,t}} z_{n,t}(j)^{\gamma-1} dj, \tag{1}$$

⁶The First Order Condition yields $Y_{n,t}^{\frac{1}{\gamma}} y_{n,t}(j)^{\frac{-1}{\gamma}} = \zeta p_{n,t}(j)$, where ζ is the Lagrange multiplier on the budget constraint. Aggregating across all j and satisfying the budget constraint yields $\zeta = Y_{n,t}^{\frac{1}{\gamma}} P_{n,t}^{\frac{1-\gamma}{\gamma}} R_t^{\frac{-1}{\gamma}}$.

and call $\tilde{z}_{n,t}(i, j) \equiv (z_{n,t}(i, j)/Z_{n,t})^{\gamma-1}$ and $\tilde{z}_{n,t}(j) \equiv (z_{n,t}(j)/Z_{n,t})^{\gamma-1}$ the “relative productivities” of firm i ’s version and of the most productive version of good j at time t , respectively. The fixed cost firm i must pay to access its idiosyncratic version of good j at time t is $\tilde{z}_{n,t}(i, j)f_S$ or $\tilde{z}_{n,t}(i, j)f_L$ if firm i is small or large, respectively.

Finally, the large firm must pay a tax τ_R multiplied by its revenue. The fixed cost and revenue tax are the only potential ex-ante differences between a small firm and the large firm in the static block. I allow for $f_S \neq f_L$ and for non-infinitesimal fixed costs so that I can lower only the large firm’s fixed cost in experiments.

2.1.3 Equilibrium

At each time t , firms choose fixed cost payments and prices to maximize static profits in a subgame perfect Nash equilibrium of the two stage game. We will see in Section 2.2.1, that for each good j , there is a unique firm with the most productive version. I focus on the equilibrium in which this is the only firm that pays the fixed cost, and the only firm with strictly positive sales of good j . Within good competition thus comes from imitations of its version, and I make the following assumption that imitation is sufficiently good so that this competition constrains its markup of price over marginal cost to be σ :

Assumption 1. $\sigma \leq \gamma/(\gamma - 1)$.

Since each good j has a unique producer at each time t , I often refer to that firm’s productivity for good j as good j ’s productivity at time t . In Appendix A, I describe equilibrium strategies, impose a restriction on fixed costs so that the most productive producer of each good finds it optimal to pay the fixed cost, and solve the static optimization problem.

2.1.4 Aggregation and Profits

Since all goods are sold with a markup σ , aggregation is straightforward. Industry output is $Y_{n,t} = Z_{n,t}L_t^P$, where L_t^P is production labor. In that sense, $Z_{n,t}$ (defined in (1)) is industry productivity. Moreover, production labor is $L_t^P = R_t/(\sigma W_t)$ and the industry price is $P_{n,t} = \sigma W_t/Z_{n,t}$, where R_t is total industry revenue and W_t is the wage.⁷

⁷The expression for L_t^P follows from setting industry revenue over variable costs equal to the cost-weighted markup, σ . The expression for $P_{n,t}$ follows from aggregating up using $p_{n,t}(j) = \sigma W_t/z_{n,t}(j)$. The expression for $Y_{n,t}$ then follows from $R_t = P_{n,t}Y_{n,t}$.

It follows that the sole producer of good j at time t earns flow profits

$$(1 - \sigma^{-1})\tilde{z}_{n,t}(j)R_t - \tilde{z}_{n,t}(j)f_S \quad (1 - \sigma^{-1} - \tau_R)\tilde{z}_{n,t}(j)R_t - \tilde{z}_{n,t}(j)f_L \quad (2)$$

if small or large, respectively. In each case, the first term is revenue minus production labor costs (and the tax for the large firm), and the second term is the fixed cost. Both terms scale linearly with the good's relative productivity, $\tilde{z}_{n,t}(j) = (z_{n,t}(j)/Z_{n,t})^{\gamma-1}$; the first due to CES demand and a constant markup across goods, and the second by definition.

2.2 Dynamic Block

2.2.1 Innovation

A firm is endowed with an ability to innovate that does not change over its lifetime, and is not a function of its past innovations. I model this ability as embedded in entrepreneurs that the firm controls. Ex-ante heterogeneity in innovation ability across firms takes the form of heterogeneity in the number of entrepreneurs firms control. I discuss the possibility of allowing a firm's number of entrepreneurs to depend on its past innovation in Section 8.

An entrepreneur has two innovation technologies. First, it can creatively destroy goods. At each time t , its firm chooses how much to target each good, described by an intensity function $\kappa : [0, M_{n,t}] \rightarrow \mathbf{R}_+$. At Poisson rate $\int_0^{M_{n,t}} \kappa(j) dj$, it creatively destroys a good, and draws the good it creatively destroys from PDF $\kappa(\cdot) / \int_0^{M_{n,t}} \kappa(j) dj$. If a firm creatively destroys good j , then it draws a new productivity for its version, $z > z_{n,t}(j)$, such that the expected value of $z^{\gamma-1}$ is $\lambda^{\gamma-1} z_{n,t}(j)^{\gamma-1}$, where $\lambda > 1$ is the innovation step size. If the firm was already producing good j , then it continues to but with a higher productivity; if not, then it becomes the producer until another firm creatively destroys or acquires good j .

Second, an entrepreneur can internally innovate. At each time t , its firm chooses a Poisson arrival rate δ at which it improves one of the firm's own goods or develops a new one. If a firm improves one of its old goods, then the productivity for its version strictly increases. If a firm develops a new good, then the set of goods expands. The firm draws a strictly positive productivity for its version of the new good, which it produces until the good is creatively destroyed or acquired by another firm. All other firms' versions start at productivity 0. It is not necessary to specify the probability that an internal innovation is a new good or an improvement on an old one because only a firm's total stock of relative

productivity across all its current goods is relevant. Conditional on an internal innovation, the expected increase in relative productivity is normalized to $\gamma - 1$.

The flow cost of an entrepreneur's innovation in units of labor is

$$\underbrace{\alpha \int_0^{M_{n,t}} \frac{\lambda^{\gamma-1} - 1}{\gamma - 1} \tilde{z}_{n,t}(j) X_S(\kappa(j)) dj + X_S(\delta)}_{\text{small firm}} \quad \underbrace{\alpha \int_0^{M_{n,t}} \frac{\lambda^{\gamma-1} - 1}{\gamma - 1} \tilde{z}_{n,t}(j) \beta X_L(\kappa(j)) dj + \beta X_L(\delta)}_{\text{large firm}}, \quad (3)$$

which depends on whether the entrepreneur is at a small firm or the large firm. In each case, the first term is the cost of creative destruction and the second term is the cost of internal innovation, where $\alpha > 0$ is a creative destruction cost shifter, $\beta > 0$ is a large firm cost shifter, and for $I \in \{S, L\}$, $X_I(\cdot)$ is a strictly increasing, strictly convex, and twice continuously differentiable cost function, with $X_I(0) = X'_I(0) = 0$. The cost of each type of innovation scales with the expected increase in relative productivity it generates divided by $\gamma - 1$ (this is the growth it generates, defined in Section 2.3). Costs are independent in the sense that the rate at which a firm creatively destroys one set of goods does not affect its cost of creatively destroying goods not in that set or of internally innovating, and vice versa. I allow for $X_S(\cdot) \neq X_L(\cdot)$ for generality, but they are equal in the quantitative model.

I allow firms to creatively destroy their own goods so that a firm's innovation technology does not depend on the set of goods it currently produces. Only a firm's stock of relative productivity is relevant, so creatively destroying one of its own goods and internally innovating have the same effect. The important distinction is between creatively destroying a competitor's good and other types of innovation: all innovation increases industry productivity $Z_{n,t}$, but if a firm creatively destroys a competitor's good, then there is also a change in the identity of the good's producer.

Innovation Cost Heterogeneity: Variation in the number of entrepreneurs is the key source of ex-ante heterogeneity across firms. A firm's cost of innovating is decreasing in its permanent number of entrepreneurs because each entrepreneur's cost is strictly convex. For simplicity, there are two types of firms: all small firms have one entrepreneur (the distribution does not matter as long as each has finitely many), and the large firm has a unit measure of entrepreneurs. Thus, the large firm's innovation rates and costs are on the same scale as the entire mass of small firms. Each small firm innovates at a Poisson arrival rate, so the entire

mass of small firms—and the large firm—innovates at a continuous rate, i.e., in finite time, it creatively destroys and internally innovates measures of goods. Thus, in equilibrium, each small firm produces finitely many goods, which are a negligible share of the total, whereas the entire mass of small firms and the large firm produce measures of goods, which are non-negligible shares of the total.

This stark distinction between each small firm and the large firm can be viewed as the limiting case of a model in which finitely many firms each have finitely many entrepreneurs, and one firm has a much larger number than the others.

2.2.2 Acquisitions

At each time t , the large firm chooses an acquisition opportunity rate function $A : [0, M_{n,t}] \rightarrow \mathbf{R}_+$. At Poisson arrival rate $A(j)$, it encounters the firm with the most productive version of good j , and makes a take-it-or-leave-it price offer subject to an acquisition tax rate $\tau_A > -1$ (a subsidy if negative). If the small firm accepts, its good is transferred and the large firm pays the offered price to the small firm as well as τ_A times the price to the tax authority. Specifically, the productivity of the large firm's version becomes the productivity of the small firm's version, and the small firm's version's productivity goes to 0. The large firm produces the good until it is creatively destroyed by another firm.

The flow labor cost of searching for acquisition opportunities is $\int_0^{M_{n,t}} \tilde{z}_{n,t}(j) X_A(A(j)) dj$, where $X_A(\cdot)$ is strictly increasing, strictly convex, and twice continuously differentiable, with $X_A(0) = X'_A(0) = 0$. As with innovation, the cost scales with the relative productivity of the good, and is independent across goods.

2.2.3 Evolution of the Measure of Small Firms

I do not yet specify the entry/exit process that drives the evolution of the measure of small firms N_t . Exit will not affect a small firm's optimal innovation (Section 4.1.4).

2.2.4 Dynamic Equilibrium

At each time t , the large firm chooses innovation rates, acquisition search rates, and acquisition price offers, and small firms choose innovation rates and acquisition price acceptance strategies to maximize expected present discounted profits using interest rate r_t . I focus on

the following Markov perfect equilibrium. The industry state is the large firm's share of industry relative productivity (or equivalently revenue, given constant markups):

$$\mathcal{L}_{n,t} \equiv \int_{j \in J_{n,t}} \tilde{z}_{n,t}(j) dj,$$

where $J_{n,t}$ is the set of goods of which the large firm has the most productive version at time t . The aggregate state is the measure of small firms N_t , industry revenue R_t , the wage W_t , and the interest rate r_t . The equilibrium is symmetric in two ways. First, a firm chooses one creative destruction rate for all its competitors' goods, one creative destruction rate for all its own goods, and one acquisition search rate for all its competitors' goods. Second, all small firms choose the same rates. Finally, in an acquisition opportunity, the large firm offers the minimum of its value of the good and the current producer's value, and the current producer accepts any offer weakly above its value.

This choice of the industry state is important because it rules out dynamic strategic behavior such as the large firm threatening to creatively destroy more intensely the goods of any small firm that creatively destroys its goods.

2.3 Growth and the Evolution of the Industry State

I now describe the dynamic equilibrium in more detail, and solve firms' optimization problems. Industry productivity grows over time as the result of innovation at rate $g_t(\mathcal{L}_{n,t}) \equiv \dot{Z}_{n,t}/Z_{n,t}$, where a dot over a variable indicates its derivative with respect to time t . Growth is the sum of growth due to small firm innovation, $N_t g_{S,t}(\mathcal{L}_{n,t})$, where each small firm generates the same expected growth, and due to large firm innovation, $g_{L,t}(\mathcal{L}_{n,t})$:

$$\begin{aligned} N_t g_{S,t}(\mathcal{L}_{n,t}) &= \frac{\lambda^{\gamma-1} - 1}{\gamma - 1} N_t \kappa_{S,t}(\mathcal{L}_{n,t}) + N_t \delta_{S,t}(\mathcal{L}_{n,t}) \\ g_{L,t}(\mathcal{L}_{n,t}) &= \frac{\lambda^{\gamma-1} - 1}{\gamma - 1} (\kappa_{L,t}(S; \mathcal{L}_{n,t})(1 - \mathcal{L}_{n,t}) + \kappa_{L,t}(L; \mathcal{L}_{n,t})\mathcal{L}_{n,t}) + \delta_{L,t}(\mathcal{L}_{n,t}). \end{aligned} \quad (4)$$

For growth due to small firm innovation, the first term is creative destruction, and the second is internal innovation. Each small firm chooses creative destruction intensity $\kappa_{S,t}(\mathcal{L}_{n,t})$ for all its competitors' goods, where the subscript S denotes that the firm is small, the subscript t indexes the aggregate state, and the argument $\mathcal{L}_{n,t}$ is the industry state in industry n at time t . The measure N_t of small firms therefore creatively destroy a representative measure

$N_t \kappa_{S,t}(\mathcal{L}_{n,t})$ of goods per unit time because every good is a competitor's good for all but one small firm out of a continuum (for the same reason, the creative destruction intensity each small firm chooses for the goods it currently produces is irrelevant). Creatively destroying a good increases its productivity raised to the $\gamma - 1$ by a factor of $\lambda^{\gamma-1} - 1$, and dividing by $\gamma - 1$ yields the growth in $Z_{n,t}$. The mass of small firms internally innovate at continuous rate $N_t \delta_{S,t}(\mathcal{L}_{n,t})$, where $\delta_{S,t}(\mathcal{L}_{n,t})$ is the rate each small firm chooses. Internal innovation generates growth in $Z_{n,t}^{\gamma-1}$ at rate $\gamma - 1$, and dividing by $\gamma - 1$ yields the growth in $Z_{n,t}$.

The expression for growth due to large firm innovation is similar (its variables are denoted by subscript L rather than S), except it does not creatively destroy all goods with the same intensity. The relative productivity weighted average intensity with which it creatively destroys goods is $\kappa_{L,t}(S; \mathcal{L}_{n,t})(1 - \mathcal{L}_{n,t}) + \kappa_{L,t}(L; \mathcal{L}_{n,t})\mathcal{L}_{n,t}$, where the argument S denotes small firm goods and L denotes large firm (its own) goods. The intensity the large firm chooses for its own goods is relevant because it conducts a non-negligible amount of innovation and produces goods with a non-negligible share of industry relative productivity.

The industry state, the large firm's share of industry relative productivity or revenue, evolves over time according to

$$\begin{aligned} \dot{\mathcal{L}}_{n,t} = & (\kappa_{L,t}(S; \mathcal{L}_{n,t}) + A_t(\mathcal{L}_{n,t}))(1 - \mathcal{L}_{n,t}) + (\gamma - 1)g_{L,t}(\mathcal{L}_{n,t})(1 - \mathcal{L}_{n,t}) \\ & - N_t \kappa_{S,t}(\mathcal{L}_{n,t})\mathcal{L}_{n,t} - (\gamma - 1)N_t g_{S,t}(\mathcal{L}_{n,t})\mathcal{L}_{n,t}. \end{aligned} \quad (5)$$

The top line is inflows. First, the large firm creatively destroys and acquires each small firm good at Poisson rates $\kappa_{L,t}(S; \mathcal{L}_{n,t})$ and $A_t(\mathcal{L}_{n,t})$, respectively, which switches the identity of the good's producer to the large firm. The total relative productivity of these goods is $1 - \mathcal{L}_{n,t}$. Second, the large firm generates growth in $Z_{n,t}^{\gamma-1}$ at continuous rate $(\gamma - 1)g_{L,t}(\mathcal{L}_{n,t})$ per unit time, which mechanically increases $\mathcal{L}_{n,t}$, its share of $Z_{n,t}^{\gamma-1}$, at this rate multiplied by $1 - \mathcal{L}_{n,t}$. Intuitively, growth depreciates the relative productivity of all other goods at the same rate, so a share $\mathcal{L}_{n,t}$ of the growth the large firm generates cannibalizes its own goods. The bottom line is outflows. Small firm creative destruction of a large firm good switches the identity of its producer to a small firm, and a share $\mathcal{L}_{n,t}$ of small firm growth comes at the expense of the relative productivity of the large firm's goods. Importantly, creatively destroying a competitor's good generates growth *and* changes the identity of the good's producer (business stealing), whereas other types of innovation only generate growth.

2.3.1 Large Firm Dynamic Optimization

I solve the large firm dynamic optimization problem, and then the small firm problem by taking the limit of an arbitrarily small large firm. I split the problem into two steps. First, subject to a given evolution of its share of industry relative productivity (the industry state), $\dot{\mathcal{L}}_{n,t}$, the large firm chooses an internal innovation rate δ , a creative destruction intensity function $\kappa(\cdot)$, and acquisition search rates $A(\cdot)$ to minimize cost. Second, it chooses $\dot{\mathcal{L}}_{n,t}$ to maximize its expected present discounted profits.

In the cost minimization problem, the large firm chooses one creative destruction intensity for all its competitors' goods, one for all its own goods, and one acquisition search rate for all its competitors' goods because the cost functions $X_L(\cdot)$ and $X_A(\cdot)$ are strictly convex, and the benefits of $\kappa(j)$ and $A(j)$ (the effect on $\dot{\mathcal{L}}_{n,t}$) and the costs scale equally with good j 's relative productivity $\tilde{z}_{n,t}(j)$. The optimal rates are determined by the first order conditions:

$$\begin{aligned}
 \text{internal innovation: } & W_t \beta X'_L(\delta_{L,t}(\mathcal{L}_{n,t})) = (\gamma - 1)(1 - \mathcal{L}_{n,t})\zeta & (6) \\
 \text{creative destruction (own): } & W_t \alpha \beta X'_L(\kappa_{L,t}(L; \mathcal{L}_{n,t})) = (\gamma - 1)(1 - \mathcal{L}_{n,t})\zeta \\
 \text{creative destruction (competitors): } & W_t \alpha \beta X'_L(\kappa_{L,t}(S; \mathcal{L}_{n,t})) = \frac{(\lambda^{\gamma-1} - 1)(1 - \mathcal{L}_{n,t}) + 1}{(\lambda^{\gamma-1} - 1)/(\gamma - 1)} \zeta \\
 \text{acquisitions: } & W_t X'_A(A_t(\mathcal{L}_{n,t})) = \max\{\zeta - (1 + \tau_A)\Pi_t(\mathcal{L}_{n,t}), 0\},
 \end{aligned}$$

where ζ is the Lagrange multiplier on the $\dot{\mathcal{L}}_{n,t}$ constraint. The left-hand sides are marginal costs and the right-hand sides are marginal benefits, relaxing the $\dot{\mathcal{L}}_{n,t}$ constraint. For innovation, both are per unit of the expected growth it generates, and for acquisitions, both are per unit of relative productivity. On the cost side, β is the large firm innovation cost shifter, α is the creative destruction cost shifter, $X_L(\cdot)$ is the large firm innovation cost function, $X_A(\cdot)$ is the acquisition search cost function, and W_t is the wage.

In all cases, the benefit comes from affecting the evolution of the large firm's share of industry relative productivity, $\dot{\mathcal{L}}_{n,t}$. The benefit is the same for internal innovation and own good creative destruction. It comes entirely from generating growth, and so is mechanically multiplied by $1 - \mathcal{L}_{n,t}$ to account for cannibalization of the relative productivity of the large firm's other goods. The benefit is higher for creatively destroying a competitor's good because it includes growth as well as the shift in the identity of the good's producer. The latter is not multiplied by $1 - \mathcal{L}_{n,t}$ because it does not affect industry productivity, and so entails no cannibalization. Finally, the benefit of an acquisition opportunity comes from shifting

the identity of the good's producer, not growth, and thus is not multiplied by $1 - \mathcal{L}_{n,t}$. It also includes the required payments to the current producer and the tax authority, where $\Pi_t(\mathcal{L}_{n,t})$ is a small firm's value of a good per unit of relative productivity.

The first order conditions for innovation hold with equality at strictly positive rates because the marginal cost at zero innovation is 0. However, the acquisition rate may be zero if the required payments exceed the benefit of taking the good. There is a unique solution given $\dot{\mathcal{L}}_{n,t}$ because each rate is increasing in ζ . Let $\bar{X}_t(\dot{\mathcal{L}}_{n,t}; \mathcal{L}_{n,t})$ be the resulting minimum cost of innovation and acquisition search/payments.⁸

The HJB equation for large firm expected discounted profits, $V_t(\mathcal{L}_{n,t})$, is

$$r_t V_t(\mathcal{L}_{n,t}) = \mathcal{L}_{n,t} \left((1 - \sigma^{-1} - \tau_R) R_t - f_L \right) + \max_{\dot{\mathcal{L}}} \left\{ \dot{\mathcal{L}} V_t'(\mathcal{L}_{n,t}) - \bar{X}_t(\dot{\mathcal{L}}; \mathcal{L}_{n,t}) \right\} + \dot{V}_t(\mathcal{L}_{n,t}), \quad (7)$$

where the right-hand side is the sum of flow profits from goods the large firm currently produces, the value of the innovation/acquisition optimization problem, and the effect of changes in the aggregate state over time. The Lagrange multiplier on the $\dot{\mathcal{L}}_{n,t}$ constraint is the marginal value of relative productivity: $\zeta = V_t'(\mathcal{L}_{n,t})$.

2.3.2 Small Firm Dynamic Optimization

A small firm's dynamic optimization problem is the same as that of the large firm except its share of industry relative productivity is negligible, and it takes the industry state as given. Thus, the analogous first order conditions for its optimal innovation rates are

$$\begin{aligned} \text{internal innovation: } & W_t X_S'(\delta_{S,t}(\mathcal{L}_{n,t})) = (\gamma - 1) \Pi_t(\mathcal{L}_{n,t}) \\ \text{creative destruction (competitors): } & W_t \alpha X_S'(\kappa_{S,t}(\mathcal{L}_{n,t})) = \frac{\lambda^{\gamma-1}}{(\lambda^{\gamma-1} - 1)/(\gamma - 1)} \Pi_t(\mathcal{L}_{n,t}). \end{aligned} \quad (8)$$

I omit creative destruction of the small firm's own goods, which is irrelevant, as discussed. On the cost side, the difference from the large firm is the small firm cost function $X_S(\cdot)$ instead of the large firm cost shifter β and function $X_L(\cdot)$.

On the benefit side, the value of generating growth is no longer multiplied by $1 - \mathcal{L}_{n,t}$ because the share that comes at the expense of the small firm's other goods is 0. Moreover, instead of the Lagrange multiplier ζ is $\Pi_t(\mathcal{L}_{n,t})$, which is the present discounted value a small

⁸ $\bar{X}_t(\dot{\mathcal{L}}_{n,t}; \mathcal{L}_{n,t}) = \beta \left[\alpha \frac{\lambda^{\gamma-1} - 1}{\gamma - 1} \left((1 - \mathcal{L}_{n,t}) X_L(\kappa_{L,t}(S; \mathcal{L}_{n,t})) + \mathcal{L}_{n,t} X_L(\kappa_{L,t}(L; \mathcal{L}_{n,t})) \right) + X_L(\delta_{L,t}(\mathcal{L}_{n,t})) \right] + (1 - \mathcal{L}_{n,t}) (X_A(A_t(\mathcal{L}_{n,t})) + A_t(\mathcal{L}_{n,t}) (1 + \tau_A) \Pi_t(\mathcal{L}_{n,t}))$.

firm gets from having the most productive version of a good per unit of the good’s current relative productivity. The value of a good is linear in its relative productivity because flow profits are linear, and all goods currently produced by small firms face the same creative destruction rate. $\Pi_t(\mathcal{L}_{n,t})$ is a function of the industry state, but not of the small firm’s past innovation, and is given by the HJB equation:

$$\begin{aligned}
r_t \Pi_t(\mathcal{L}_{n,t}) &= (1 - \sigma^{-1}) R_t - f_S + \dot{\mathcal{L}}_{n,t} \Pi'_t(\mathcal{L}_{n,t}) + \dot{\Pi}_t(\mathcal{L}_{n,t}) \\
&\quad - \underbrace{\left(N_t \kappa_{S,t}(\mathcal{L}_{n,t}) + \kappa_{L,t}(S; \mathcal{L}_{n,t}) \right)}_{\text{creative destruction risk}} + \underbrace{(\gamma - 1) g_t(\mathcal{L}_{n,t})}_{\text{growth}} \Pi_t(\mathcal{L}_{n,t}).
\end{aligned} \tag{9}$$

small firm depreciation rate

The right-hand side of the first line is flow profits and the change in the value over time due to changes in the industry and aggregate state. The second line is the effect of creative destruction risk and growth on the value of a good.

The Small Firm Depreciation Rate: Equation (9) shows that the effective discount rate a small firm applies to an innovation is the interest rate, r_t , plus the “small firm depreciation rate”, which is the sum of creative destruction risk and the growth rate of $Z_{n,t}^{\gamma-1}$. At a Poisson arrival rate, the identity of the good’s producer shifts to another firm (creative destruction risk), and at a continuous rate, industry growth erodes the good’s relative productivity.

3 Key Mechanisms and Qualitative Results

I first decompose the growth rate into the small firm depreciation rate, and growth relative to small firm depreciation. The latter reflects the direction of innovation and is exogenous in standard growth models. Here, it depends on the large firm’s contribution to small firm depreciation and share of industry revenue, which drives the results that follow.

I then characterize the long-run effects of two types of permanent shocks. The first is a shock that only directly affects the large firm, which I call a shock to “large firm innovation incentives”. This is the type of shock I use in the quantitative experiment in Section 5 to explain the recent fall in the long-run growth rate. Moreover, it illustrates the effects of a size-dependent tax. I discuss how the results differ for a shock that directly affects small firms. The second shock I study is to the acquisition tax rate τ_A .

I begin with a special case of the model in which firms only find it worthwhile to innovate by creatively destroying their competitors' goods, and then return to the full model. In the full model, two separate mechanisms drive the crucial relationship between the large firm's size and the growth rate relative to the small firm depreciation rate, whereas in the special case, only one is present. Nonetheless, the implications of this relationship for the main results are the same. The special case is simpler, and therefore useful for illustrating these implications. Moreover, the special case and the full model together demonstrate that the main results are robust to whether or not firms internally innovate.

3.1 Only Creative Destruction of Competitors' Goods

For this subsection, I simplify the model by setting $\gamma = 1$. For the model to be properly specified without a love of variety, I fix the measure of goods, $M_{n,t}$, to be 1. To map the creative destruction process into this setting, suppose that when a firm creatively destroys good j , they draw a new productivity for their version, $z > z_{n,t}(j)$, such that the expected value of z is $\lambda z_{n,t}(j)$, where $\lambda > 1$ is the innovation step size.

The key implication of $\gamma = 1$ is that all goods have relative productivity 1, and so receive revenue R_t and the same flow profits, regardless of productivity (conditional on whether the producer is small or large).⁹ It follows that firms do not benefit from creatively destroying their own goods or internally innovating.¹⁰ Moreover, the value to a firm of having the most productive version of a good depreciates over time due only to creative destruction risk, not growth.¹¹ Finally, there is a Markov perfect equilibrium without an industry state; the large firm's creative destruction intensity for its competitors' goods does not depend on its share of industry revenue because it no longer has a cannibalization effect on its own goods.¹²

Since firms do not internally innovate, I ignore it and set the relative cost of creative destruction, α , to 1. I use the same notation for innovation/acquisition rates and value functions as before, but drop the $\mathcal{L}_{n,t}$ argument because it is no longer the industry state. I also drop the S vs. L argument for the large firm innovation rate $\kappa_{L,t}$ because the large firm only targets small firm goods.

⁹See (2) for flow profits.

¹⁰See the first order conditions for $\kappa_{L,t}(L; \mathcal{L}_{n,t})$ and $\delta_{L,t}(\mathcal{L}_{n,t})$ in (6), where the marginal benefit is now 0.

¹¹See (9) for small firms, and (5) for the large firm.

¹²See the first order condition for $\kappa_{L,t}(S; \mathcal{L}_{n,t})$ in (6), where the benefit no longer depends on $\mathcal{L}_{n,t}$.

In this Markov perfect equilibrium, the large firm’s share of industry revenue, or equivalently of industry goods, $\mathcal{L}_{n,t}$, evolves over time according to

$$\dot{\mathcal{L}}_{n,t} = (\kappa_{L,t} + A_t)(1 - \mathcal{L}_{n,t}) - N_t\kappa_{S,t}\mathcal{L}_{n,t}.$$

Inflows are the large firm creatively destroying and acquiring each of the measure $1 - \mathcal{L}_{n,t}$ of its competitors’ goods at Poisson rates $\kappa_{L,t}$ and A_t , respectively, shifting the identity of the good’s producer to the large firm. Outflows are the mass of small firms creatively destroying each of the measure $\mathcal{L}_{n,t}$ of the large firm’s goods at Poisson rate $N_t\kappa_{S,t}$.

3.1.1 Industry Concentration and Growth

To see the key mechanism in this special case of the model, decompose growth as

$$g_t(\mathcal{L}_{n,t}) \equiv \frac{\dot{Z}_{n,t}}{Z_{n,t}} = (\kappa_{L,t} + N_t\kappa_{S,t}) \left(\frac{N_t\kappa_{S,t}}{\kappa_{L,t} + N_t\kappa_{S,t}} \ln(\lambda) + \frac{\kappa_{L,t}}{\kappa_{L,t} + N_t\kappa_{S,t}} (1 - \mathcal{L}_{n,t}) \ln(\lambda) \right), \quad (10)$$

where industry productivity is $Z_{n,t} = \exp\left(\int_0^1 \ln(z_{n,t}(j)) dj\right)$, and growth depends on the large firm’s share of industry revenue even though each firm’s innovation decisions do not. The first term is the small firm depreciation rate. In this special case of the model, it is the rate at which a small firm’s good is creatively destroyed—small firm creative destruction risk. This determines the time a small firm expects to earn profits from an innovation, which is the channel through which innovation by other firms affects a small firm.

The second term is the growth rate relative to the small firm depreciation rate, which is the average of the growth each firm generates relative to the depreciation it imposes on small firms, where each firm is weighted by its contribution to small firm depreciation. A small firm’s innovation targets all goods, so its ratio is $\ln(\lambda)$, the growth generated by an innovation. Small firms as a whole receive weight $N\kappa_{S,t}/(N\kappa_{S,t} + \kappa_{L,t})$. The large firm’s innovation targets only small firm goods, so its ratio is $(1 - \mathcal{L}_{n,t})\ln(\lambda)$, which reflects that not targeting the measure $\mathcal{L}_{n,t}$ of goods it produces reduces the growth the large firm generates, but does not reduce the depreciation it imposes on small firms. I call this the “Composition Effect” of large firm innovation because it follows mechanically from the distribution of goods across firms. It implies that the overall growth rate relative to the small firm depreciation rate is falling in the large firm’s shares of small firm depreciation and industry revenue.

We can easily extend the decomposition for any number of large firms, each of which targets all its competitors’ goods with the same creative destruction intensity, but not its

own. If there are no acquisitions, then long-run growth relative to small firm depreciation is $\ln(\lambda)(1 - HHI)$, where HHI is the Herfindahl–Hirschman index for industry revenue.

To compare to the same growth rate decomposition in a standard creative destruction model without a large firm, such as Klette and Kortum (2004), set $\kappa_{L,t} = 0$. The difference is that growth relative to small firm depreciation is $\ln(\lambda)$, which does not depend on equilibrium outcomes. Indeed, this ratio is exogenous in most endogenous growth models with creative destruction and/or expanding varieties, as long as all firms face the same depreciation rate due to innovation.¹³ Here in an oligopolistic growth model, it is falling in the concentration of innovation and revenue. Large firm innovation targets their competitors’ goods, so the growth they generate is low relative to the depreciation they impose on their competitors.

For what follows, I study balanced growth path equilibria in which industry revenue, the wage, the interest rate, and the measure of small firms are constant over time at R , W , r , and N , and the large firm’s share of industry revenue is constant over time at $\mathcal{L} = (\kappa_L + A)/(\kappa_L + A + N\kappa_S)$. It follows that the growth rate is constant over time as well.

3.1.2 Shocks to Large Firm Innovation Incentives

I first set acquisitions to 0 and study the long-run effects of permanent shocks and policy changes that only directly affect large firm innovation incentives. Specifically, I vary the large firm fixed cost f_L , innovation cost shifter β , and revenue tax rate τ_R , and characterize the implied long-run relationship between the rate at which the large firm creatively destroys each of its competitors’ goods, κ_L , the large firm’s share of industry revenue, \mathcal{L} , and the growth rate, g . In Appendix C, I provide a mild condition under which there is a unique equilibrium; in that case, equilibrium κ_L is a continuous increasing function of flow profits relative to β , $((1 - \sigma^{-1} - \tau_R)R - f_L)/\beta$, and does not depend otherwise on f_L , β , or τ_R . We can therefore use the results to back out exactly the long-run effects of permanent shocks and policy changes on concentration and growth.

I begin with the following theorem, which characterizes whether a marginal increase in the large firm’s innovation incentives is associated with faster growth in the long-run.

¹³For example, it holds in Aghion, Bergeaud, Boppart, Klenow, and Li (2022) with process productivity heterogeneity, and in Atkeson and Burstein (2019) with creative destruction and expanding varieties. An exception is De Ridder (2021) in which more productive firms face a lower creative destruction rate.

Theorem 1. *Let $\partial g/\partial\kappa_L$ be the change in the long-run growth rate relative to the change in the equilibrium rate at which the large firm creatively destroys each of its competitors' goods, following a marginal change in the large firm fixed cost f_L , innovation cost shifter β , or revenue tax rate τ_R . Its sign—whether it is positive or negative—is given by*

$$\text{sign}\left(\frac{\partial g}{\partial\kappa_L}\right) = \text{sign}\left(\frac{1 - \mathcal{L}r + \kappa_L + N\kappa_S}{2\mathcal{L}\kappa_L + N\kappa_S} - \epsilon_S\right), \quad (11)$$

where ϵ_S is the elasticity of a small firm's innovation, κ_S , with respect to the small firm's value of an innovation, Π .

Proof. See Appendix B. ■

Following a rise in the rate at which the large firm creatively destroys each small firm good, κ_L , small firms respond by innovating less, and there are two counteracting effects on growth, each corresponding to one of the terms in the growth rate decomposition in (10). First, the total equilibrium rate at which a small firm good is creatively destroyed increases. As small firms innovate less, their marginal cost of innovating falls (costs are convex), so their first order condition is satisfied with a higher creative destruction rate.

Second, the “Composition Effect” implies a fall in growth relative to the rate at which a small firm good is creatively destroyed. The large firm's share of small firm creative destruction risk rises, which increases the weight large firm innovation receives in (10). Moreover, as a result, the large firm's long-run industry revenue share increases, which reduces the growth the large firm generates relative to the creative destruction it imposes on small firm goods.

Theorem 1 states that the second effect dominates—growth falls—if three observable equilibrium outcomes are sufficiently high: the elasticity of small firm innovation; the creative destruction risk share of the discount rate on a small firm's good, which also includes the interest rate; and the large firm's shares of revenue and small firm creative destruction risk, which are equal in the long-run without acquisitions. The first two imply that an increase in large firm innovation leads to a big fall in small firm innovation, so the total rate at which a small firm good is creatively destroyed increases by little, but concentration increases by a lot. The third makes a rise in concentration worse for growth because the large firm's shares of small firm creative destruction risk and revenue strengthen each others' effects.

The following corollary concerns the effects of non-marginal changes in large firm innovation incentives. If those incentives are the main driver of variation across industries or

over time, then it characterizes the long-run relationship between industry concentration and growth, and between industry concentration and the growth effects of a tax on large firms.

Corollary 1. *Vary the large firm fixed cost f_L , innovation cost shifter β , and revenue tax rate τ_R . The long-run large firm revenue share \mathcal{L} is a strictly increasing function of the rate at which the large firm creatively destroys each small firm good, κ_L , and the long-run growth rate g is a continuously differentiable function of κ_L such that*

1. for a “small” large firm: $\lim_{\kappa_L \rightarrow 0} (\partial g / \partial \kappa_L) > 0$;
2. for a dominant large firm: $\lim_{\kappa_L \rightarrow \infty} \left(\frac{\partial g / \partial \kappa_L}{g} \right) < 0$ and $\lim_{\kappa_L \rightarrow \infty} (g) = 0$;
3. and if ϵ_S is constant, then g is single-peaked, and so increasing then decreasing.

Proof. See Appendix B. ■

First, a “small” large firm targets almost all goods with creative destruction like a small firm, so if it innovates more, then the “Composition Effect” is absent, and growth increases. Second, a dominant large firm creatively destroys its competitors’ goods quickly, but generates little growth because in the long-run, there are few small firm goods to creatively destroy. Third, if the elasticity of small firm innovation with respect to the value of innovating is constant (due to an isoelastic cost function), then the first two features are the whole story: growth exhibits an inverted-U shape as a function of the large firm’s industry revenue share. In terms of policy, the corollary implies that a tax on large firms is more likely to increase long-run growth in more concentrated industries.

Finally, for completeness and to illustrate the importance of concentration, I compute the expression analogous to (11) in Theorem 1, for a shock to *small firm* innovation incentives:¹⁴

$$\text{sign} \left(\frac{\partial g}{\partial N \kappa_S} \right) = \text{sign} \left(\frac{1 + \mathcal{L}^2}{\mathcal{L}(1 - \mathcal{L})^2} \frac{r + \kappa_L + N \kappa_S}{\kappa_L + N \kappa_S} - \epsilon_L \right),$$

where ϵ_L is the *elasticity of the large firm innovation rate, κ_L , with respect to the large firm’s value of an innovation, V'* . The difference from (11) is that $(1 - \mathcal{L})/2$, which is decreasing in \mathcal{L} from 1/2 to 0, is replaced by $(1 + \mathcal{L}^2)/(1 - \mathcal{L})^2$, which is increasing from 1 to infinity. Thus, an increase in small firm innovation can reduce growth if it engenders a sufficient large

¹⁴This follows from fully differentiating the large firm value of a good, $V' = \frac{(1 - \sigma^{-1} - \tau_R)R - f_L + W \ln(\lambda) \beta X_L(\kappa_L)}{r + \kappa_L + N \kappa_S}$, with respect to an exogenous change in κ_S , and using that the elasticity of κ_L with respect to V' is ϵ_L .

firm response, but the required large firm elasticity, ϵ_L , is much higher than the required small firm elasticity, ϵ_S , for an increase in large firm innovation to reduce growth.

Appendix C shows that in a stable (or unique) equilibrium, either $\partial g/\partial \kappa_L \geq 0$ or $\partial g/\partial N\kappa_S \geq 0$. In that case, for an increase in small firm innovation to reduce growth, the elasticity of large firm innovation, ϵ_L , must be much higher than the elasticity of small firm innovation, ϵ_S . For example, twice as high if the large firm's share of industry revenue is near 0, and more than 11 times as high if it is 40.7% as in the calibrated model. On the other hand, an increase in large firm innovation can reduce growth even if small and large firm innovation are equally elastic.

3.1.3 Acquisition Policy

I now allow for acquisitions, and study the long-run effects of changes to the acquisition tax rate, τ_A . I restrict attention to the limiting case in which the elasticity of small firm innovation with respect to the value of innovating, ϵ_S , is approaching infinity (the cost function is approaching linearity). This is relevant for analyzing policy in the macroeconomic model in which a free entry condition makes small firm innovation perfectly elastic. This restriction holds fixed the rate at which a small firm's good is creatively destroyed, so an increase in large firm innovation incentives reduces growth. I impose the condition in Appendix C for a unique equilibrium, which is $\epsilon_L < 1$ given $\epsilon_S \rightarrow \infty$. The following theorem characterizes whether a marginal increase in the acquisition tax rate leads to faster growth in the long-run.

Theorem 2. *Suppose the acquisition rate, A , is strictly positive, and the small firm innovation cost function, $X_S(\cdot)$, is approximately linear. The sign of the derivative of the long-run growth rate with respect to the acquisition tax rate is given by*

$$\text{sign}\left(\frac{\partial g}{\partial \tau_A}\right) = \text{sign}\left(\epsilon_A - \frac{1}{1 - \mathcal{L}} \frac{2\kappa_L + A}{(r + \kappa_L + N\kappa_S)/\epsilon_L - \kappa_L} \left(1 + \frac{(1 + \tau_A)\Pi}{V' - (1 + \tau_A)\Pi}\right)^{-1}\right), \quad (12)$$

where ϵ_L is the elasticity of the large firm innovation rate, κ_L , with respect to the value of an innovation, V' , and ϵ_A is the elasticity of the acquisition rate, A , with respect to the surplus of each acquisition, $V' - (1 + \tau_A)\Pi$.

Proof. See Appendix D. ■

Following a rise in the acquisition tax rate, there are two counteracting effects on growth. First, the rate at which the large firm acquires each small firm good, A , falls. This reduces the

concentration of revenue, which increases the rate at which the large firm generates growth relative to the creative destruction it imposes on small firm goods. As a result, growth rises. This is the effect an econometrician would estimate from an exogenous acquisition.

Second, the rate at which the large firm creatively destroys each competitor's good, κ_L , falls. The large firm internalizes that creatively destroying a competitor's good reduces the set of goods available to acquire, so a fall in the value of acquiring raises the incentive to innovate. Thus, the concentration of innovation increases, which leads to lower growth. This effect is not present following a single exogenous acquisition because it depends on the large firm's expectation of future valuable acquisitions.

Theorem 2 states which effect dominates as a function of observable equilibrium outcomes. In particular, increasing the acquisition tax rate reduces growth if large firm innovation is sufficiently responsive, if the acquisition rate is sufficiently unresponsive, and if the acquisition surplus is sufficiently high. Then, the acquisition rate does not fall by much, but the surplus of each acquisition does, which leads to a big rise in large firm innovation.

The following corollary, similar to Corollary 1, supposes that variation in large firm innovation incentives drives variation across industries or over time, and characterizes the relationship between concentration and the effect of a change in the acquisition tax rate. For any characterization to be possible, I must impose restrictions on the elasticities of large firm innovation and the acquisition rate with respect to the value of innovating and of acquiring, respectively, because they are important for the effects of changes in the acquisition tax rate.

Corollary 2. *Suppose the acquisition rate, A , is strictly positive, the large firm innovation and acquisition search cost functions, $X_L(\cdot)$ and $X_A(\cdot)$, are isoelastic, and the small firm innovation cost function, $X_S(\cdot)$, is approximately linear. Vary the large firm fixed cost, f_L , and revenue tax rate, τ_R . There exists an $\mathcal{L}^* \in (0, 1)$ such that the derivative of the long-run growth rate with respect to the acquisition tax rate, $\partial g / \partial \tau_A$, is strictly negative if and only if the long-run large firm share of industry revenue, \mathcal{L} , is strictly greater than \mathcal{L}^* .*

Proof. See Appendix D. ■

Taxing acquisitions is worse for growth in more concentrated industries, and following a general rise in concentration driven by large firm innovation incentives, as this paper argues occurred recently in the US. Intuitively, long-run growth is decreasing in the product of the large firm's shares of industry revenue and small firm creative destruction risk. As these

shares go to 1, mechanically, the partial derivative of the product with respect to the small firm innovation rate becomes infinitely bigger than the partial derivatives with respect to the rates at which the large firm creatively destroys and acquires each of its competitors' goods. Therefore, taxing acquisitions reduces growth because it encourages large firm innovation, which discourages small firm innovation.

3.2 All Innovation Types

I now set the elasticity of substitution across goods, γ , to be strictly greater than 1, and return to the full model of Section 2. As such, the revenue and profits a good yields is linear and increasing in its relative productivity, $(z_{n,t}(j)/Z_{n,t})^{\gamma-1}$. Firms therefore creatively destroy their own goods as well as their competitors', and internally innovate. Moreover, the value to a firm of an innovation depreciates over time due to creative destruction risk and the continuous erosion of relative productivity by growth. Specifically, the small firm depreciation rate is $\kappa_{L,t}(S; \mathcal{L}_{n,t}) + N_t \kappa_{S,t}(\mathcal{L}_{n,t}) + (\gamma - 1)g_t(\mathcal{L}_{n,t})$. The industry state is $\mathcal{L}_{n,t}$, the large firm's share of industry relative productivity or, equivalently, revenue.

The difference from the analysis of the $\gamma = 1$ case is in the impact of the large firm's industry revenue share on the rate at which the large firm generates growth relative to the depreciation it imposes on its competitors. As before, this ratio depends on the growth the large firm generates relative to the creative destruction risk it imposes on small firms,¹⁵ which is now driven by the "Composition Effect", and a new "Relative Innovation Effect".

I compute the relative rates firms choose for different types of innovation, and show how they determine the "Composition Effect" and "Relative Innovation Effect". I then characterize when the effects of changes in large firm innovation incentives are as before.

I first impose a functional form on the innovation cost functions $X_S(\cdot)$ and $X_L(\cdot)$, which holds for the rest of the paper. Since a firm chooses multiple innovation rates, their relative rates are crucial. I thus make $X_S(\cdot)$ and $X_L(\cdot)$ isoelastic, with the same elasticity:

$$\text{for all } x \geq 0, \quad X_S(x) = X_L(x) = (1/\epsilon + 1)^{-1} x^{1/\epsilon+1}, \quad (13)$$

where $\epsilon > 0$. It follows that for all firms and all innovation types, the elasticity of the

¹⁵The ratio is now given by $g_{L,t}(\mathcal{L}_{n,t})/(\kappa_{L,t}(S; \mathcal{L}_{n,t}) + (\gamma - 1)g_{L,t}(\mathcal{L}_{n,t}))$, which is a strictly increasing *non-linear* function of $g_{L,t}(\mathcal{L}_{n,t})/\kappa_{L,t}(S; \mathcal{L}_{n,t})$.

innovation rate with respect to the value of that type of innovation to that firm is ϵ .¹⁶

From the first order conditions in (6), the large firm's relative innovation rates are:

$$\begin{aligned} \frac{\text{creative destruction (own)}}{\text{creative destruction (competitors)}}: \quad & \frac{\kappa_{L,t}(L; \mathcal{L}_{n,t})}{\kappa_{L,t}(S; \mathcal{L}_{n,t})} = \left(\frac{1 - \mathcal{L}_{n,t}}{1 - \mathcal{L}_{n,t} + 1/(\lambda^{\gamma-1} - 1)} \right)^\epsilon \\ \frac{\text{internal innovation}}{\text{creative destruction (competitors)}}: \quad & \frac{\delta_{L,t}(\mathcal{L}_{n,t})}{\kappa_{L,t}(S; \mathcal{L}_{n,t})} = \left(\alpha \frac{1 - \mathcal{L}_{n,t}}{1 - \mathcal{L}_{n,t} + 1/(\lambda^{\gamma-1} - 1)} \right)^\epsilon, \end{aligned} \quad (14)$$

where the first is the rate at which it creatively destroys each of its own goods relative to the rate for each of its competitors' goods, and the second is the rate at which it internally innovates relative to the same. Each relative rate reflects that the component of an innovation's relative productivity from generating growth is discounted by $1 - \mathcal{L}_{n,t}$ due to cannibalization, whereas the component from taking a competitor's good is not. The latter is only present for creatively destroying a competitor's good, and the relative productivity it entails over the relative productivity from growth is $1/(\lambda^{\gamma-1} - 1)$. The only difference between the two relative rates is the relative cost of creative destruction, α . Finally, a small firm's relative innovation rates are given by setting $\mathcal{L}_{n,t} = 0$.

There are two important features of the relative innovation rates. First, a firm chooses a lower creative destruction rate for each of its own goods than for its competitors' (the first relative rate is less than 1) because the latter yields additional business stealing. Thus, the "Composition Effect" is still present: holding fixed relative rates, the rate at which the large firm generates growth *through creative destruction* relative to the rate at which it creatively destroys each small firm good is falling in its share of industry revenue.

Second, each relative innovation rate is decreasing in the large firm's industry revenue share, $\mathcal{L}_{n,t}$, because a higher share reduces the value of generating growth, but not of taking a competitor's good. Thus, the rate at which the large firm generates growth *through all types of innovation* relative to the rate at which it creatively destroys each small firm good is falling in its industry revenue share. I call this the "Relative Innovation Effect". It relates to the Arrow (1962) replacement effect that incumbents are reluctant to innovate for fear of cannibalization. Here, cannibalization affects some types of innovation more than others.

Using the calibrated model, Figure 1 displays, as functions of the large firm's industry revenue share, the large firm's equilibrium relative innovation rates, the rate at which the large firm generates growth relative to the depreciation it imposes on small firms, and the

¹⁶The first order conditions ((6) for the large firm and (8) for small firms) show that each innovation rate is proportional to the value of that type of innovation to the power of ϵ .

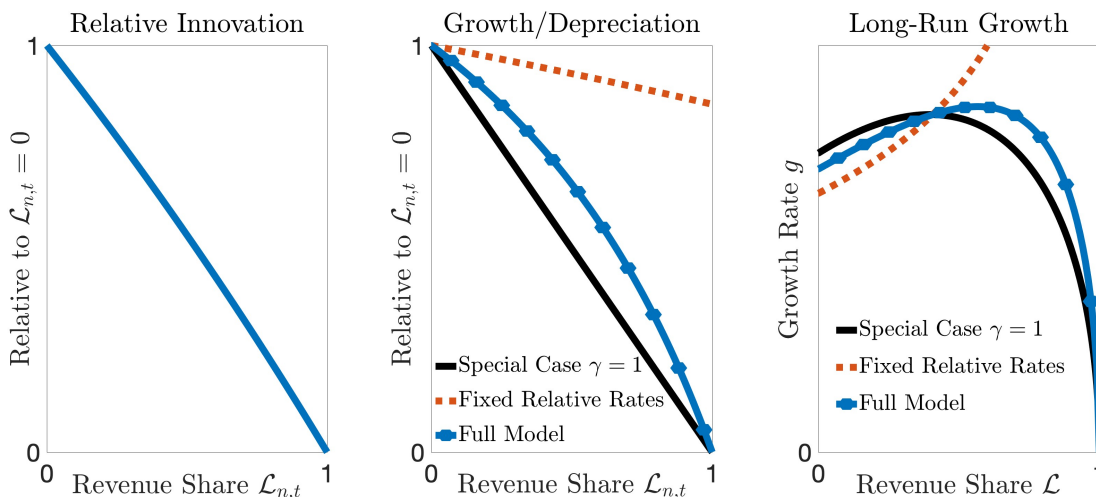


Figure 1: Left panel: large firm relative innovation rates from (14). Middle panel: large firm growth over the depreciation it imposes on small firms, under three scenarios: the special case from Section 3.1 with $\gamma = 1$, the full model with relative innovation fixed at $\mathcal{L}_{n,t} = 0$, and the full model. Values are relative to $\mathcal{L}_{n,t} = 0$ (the small firm value). Right panel: long-run growth due to shocks to large firm innovation incentives under the three scenarios. Full model figures use the initial calibrated model, and the special case calibration is discussed in Section 5.2.

long-run industry growth rate following shocks to large firm innovation incentives. The latter two are under three scenarios: the special case with $\gamma = 1$ and so only creative destruction of competitors' goods; with $\gamma > 1$ but relative innovation rates fixed at their $\mathcal{L}_{n,t} = 0$ values; and the full model with equilibrium relative innovation rates.

The three scenarios illustrate a concern with creative destruction models, and show that it does not apply here. Namely, that mechanisms related to creative destruction will be insignificant if we introduce internal innovation, which is a more prominent source of growth in the data. In the first two scenarios, only the ‘‘Composition Effect’’ is present. In the second, it is less significant because it only affects creative destruction, and firms also internally innovate. As a result, growth goes to infinity rather than 0 as the large firm becomes a monopolist because creative destruction goes to 0, but the large firm continues to internally innovate. However in the full model, the ‘‘Relative Innovation Effect’’ compensates, and the inverted-U returns. The following theorem, comparable to Corollary 1, shows that the ‘‘Relative Innovation Effect’’ compensates if and only if the elasticity ϵ is sufficiently high.

Theorem 3. *Vary the large firm fixed cost f_L , innovation cost shifter β , and revenue tax*

rate τ_R . The long-run large firm revenue share \mathcal{L} is a strictly increasing function of the long-run rate at which the large firm creatively destroys each small firm good, $\kappa_L(S)$, and the long-run growth rate g is a continuously differentiable function of $\kappa_L(S)$ such that

1. for a “small” large firm: $\lim_{\kappa_L(S) \rightarrow 0} (\partial g / \partial \kappa_L(S)) > 0$;
2. for a dominant large firm if $\epsilon > \frac{\sqrt{5}-1}{2}$: $\lim_{\kappa_L(S) \rightarrow \infty} \left(\frac{\partial g / \partial \kappa_L(S)}{g} \right) < 0$ and $\lim_{\kappa_L(S) \rightarrow \infty} (g) = 0$;
3. for a dominant large firm if $\epsilon < \frac{\sqrt{5}-1}{2}$: $\lim_{\kappa_L(S) \rightarrow \infty} \left(\frac{\partial g / \partial \kappa_L(S)}{g} \right) > 0$ and $\lim_{\kappa_L(S) \rightarrow \infty} (g) = \infty$;
4. and if $\epsilon = 1$, then g is single-peaked, and so increasing then decreasing.

Proof. See Appendix E. ■

The difference from Corollary 1 is that if the elasticity ϵ is sufficiently low, then growth is increasing in large firm innovation incentives even at high levels of concentration. Whereas if $\epsilon = 1$, then growth again exhibits an inverted-U shape as a function of the large firm’s industry revenue share. Intuitively, following a rise in its innovation incentives, the large firm innovates more, and the “Relative Innovation Effect” shifts its innovation toward creatively destroying its competitors’ goods. For growth to fall, the second force must dominate, leading to a fall in large firm internal innovation. This holds only if ϵ is sufficiently high.

4 Quantifying the Model

I now quantify the model to analyze the relevance of the mechanisms for the recent rise in industry concentration and fall in growth in the US, and to see the implications for welfare and policy. It is important first to embed the industry model of Section 2 into a macroeconomic model to allow for general equilibrium effects and compute welfare.

4.1 The Macroeconomy

There is a unit measure of ex-ante identical industries, indexed by n . There is a single large firm in each industry, which takes macroeconomic aggregates as given. A representative final good producer aggregates industry goods to sell to a representative household for consumption and to firms for fixed costs. The household consumes, inelastically supplies labor, owns

all firms, and has access to a risk-free bond in zero supply. Exogenous exit and undirected endogenous entry determine the measure of small firms, which is the same in each industry.

4.1.1 The Household and the Interest Rate

The household chooses a path of consumption, $\{C_t\}$, and bond holdings, $\{B_t\}$, to maximize present discounted utility subject to the budget constraint at all t :

$$\max \int_0^{\infty} e^{-\rho t} \ln(C_t) dt \quad \text{s.t.} \quad P_t C_t + \dot{B}_t = W_t \bar{L} + D_t + r_t B_t \quad \text{for all } t,$$

taking as given paths for the final good price, $\{P_t\}$, wage, $\{W_t\}$, firm profits, $\{D_t\}$, and interest rate, $\{r_t\}$, where $\rho > 0$ is the time discount rate and \bar{L} is exogenous labor supply.

Bond market clearing implies that for all t , the net interest rate must compensate for ρ and the negative rate of change of the value of the numeraire: $r_t = \rho + \dot{C}_t/C_t + \dot{P}_t/P_t$.

4.1.2 The Final Good Producer and Aggregate Expenditure

The final good producer aggregates industry goods, $\{Y_{n,t}\}$, into final good output, Y_t , via a Cobb-Douglas aggregator: $\ln(Y_t) = \int_0^1 \ln(Y_{n,t}) dn$. It maximizes profits taking as given the final good price, P_t , and industry prices, $\{P_{n,t}\}$. The First Order Condition for $Y_{n,t}$ implies constant expenditures across industries: $R_t = P_t Y_t$. Perfect competition implies zero profits, which yields the final good price: $\ln(P_t) = \int_0^1 \ln(P_{n,t}) dn$.

4.1.3 Aggregation and Normalization

Since each industry is identical, final good and aggregate quantities and prices are equal to their industry counterparts, derived in Section 2.1.4. I normalize the final good price to be $P_t = Z_t^{-1}$ for all t , where Z_t is aggregate productivity, so the wage is $W_t = \sigma^{-1}$, and aggregate expenditure is $R_t = L_t^P$. I impose that a unit of fixed cost requires Z_t units of the final good so that with the normalization, fixed costs are as before (Section 2.1.2).

4.1.4 Entry and Exit

At each time t , each of an infinite mass of potential entrants can receive value 0 or pay $\xi > 0$ units of labor to draw an industry from the uniform distribution and enter as a small

firm with one entrepreneur (any finite distribution is the same) and a 0 productivity version of each good. Each small firm entrepreneur exits exogenously at Poisson rate $\eta > 0$, after which it cannot innovate; its firm survives until all its goods are creatively destroyed. Thus, the measure of small firm *entrepreneurs*, which is relevant, does not equal the measure of small firms, which is irrelevant. The former, N_t , is the same in each industry because entry is *undirected*, and evolves over time due to entry e_t and exit according to $\dot{N}_t = e_t - \eta N_t$.

Potential entrants maximize expected discounted profits using the interest rate r_t . The value of entering industry n at time t is $E_t(\mathcal{L}_{n,t})$, given by the HJB equation:

$$\begin{aligned} r_t E_t(\mathcal{L}_{n,t}) = & W_t (1 + \epsilon)^{-1} \left(\alpha \frac{\lambda^{\gamma-1} - 1}{\gamma - 1} \kappa_{S,t} (\mathcal{L}_{n,t})^{1/\epsilon+1} + \delta_{S,t} (\mathcal{L}_{n,t})^{1/\epsilon+1} \right) \\ & - \eta E_t(\mathcal{L}_{n,t}) + \dot{\mathcal{L}}_{n,t} E'_t(\mathcal{L}_{n,t}) + \dot{E}_t(\mathcal{L}_{n,t}), \end{aligned} \quad (15)$$

where the first line is the flow value of innovating, which includes the expected discounted profits from any innovations.¹⁷ The second line is the effect of exit, in which case the flow value of innovating goes to 0, and the effect of changes in the industry and aggregate states over time. The value of entry net the cost is $\int_0^1 E_{n,t}(\mathcal{L}_{n,t}) dn - W_t \xi$, which must be weakly negative for the entry rate to be finite.

4.1.5 Macroeconomic Equilibrium

The household, final good producer, and potential entrants optimize, taking as given paths of the state, $\{N_t\}$, and aggregates, $\{R_t, W_t, r_t\}$. Final good producers earn zero static profits, the value of entry is weakly less than the cost, the tax authority uses lump sum taxes on the household to balance its budget, and markets clear for bonds, labor (production, innovation, and entry equal supply), and the final good (consumption and fixed costs equal supply).

I study balanced growth paths and convergence to them following unanticipated shocks. A balanced growth path is an equilibrium in which each industry state and the aggregate state are constant over time. Thus, aggregate quantities and prices, and industry and aggregate growth rates are constant, and the interest rate is $r = \rho$.

¹⁷The flow value uses the first order conditions in (8), and the innovation cost in (3) with function (13).

4.1.6 Welfare

The measure of welfare is that of the household, which depends on current productivity and expected discounted future growth, production labor, and fixed costs:

$$\int_0^{\infty} e^{-\rho t} \ln(C_t) dt = \frac{\ln(Z_0)}{\rho} + \int_0^{\infty} e^{-\rho t} \frac{g_t}{\rho} dt + \int_0^{\infty} e^{-\rho t} \ln \left(L_t^P - \int_0^1 ((1 - \mathcal{L}_{n,t}) f_S + \mathcal{L}_{n,t} f_L) dn \right) dt.$$

The first two terms are expected discounted log productivity, where $g_t = \dot{Z}_t/Z_t$ is the aggregate productivity growth rate, and the last term is expected discounted log consumption relative to productivity because consumption is final good output minus fixed costs. Growth is divided by ρ because growth at time t raises consumption permanently, and so is equivalent to a $1/\rho$ increase in log consumption at time t .

4.2 Calibration

I calibrate an initial balanced growth path. I set some parameters externally (Table 1) and internally calibrate the rest to match empirical moments (Table 2). I set labor supply, \bar{L} , so aggregate expenditure, $R = L^P$, is 1. Small and large firms have the same fixed cost, $f_S = f_L$. Innovation cost functions, $X_S(\cdot)$ and $X_L(\cdot)$, are identical and isoelastic as in (13). The revenue tax rate is 0. There are no acquisitions. The units of time are years.

Externally Calibrated Parameters: The exit rate, η , is the annual employment-weighted average firm exit rate from Boar and Midrigan (2022).¹⁸ The demand elasticity, γ , is the median estimate from Broda and Weinstein (2006) at the most disaggregated level in the earliest time period, which fits with using the most disaggregated industry definition for the large firm revenue share measure. The inverse marginal innovation cost elasticity, ϵ , follows a variety of studies discussed in Acemoglu, Akcigit, Alp, Bloom, and Kerr (2018). Moreover, it is near the lower bound below which the model does not generate the inverted-U relationship between growth and concentration documented in Cavenaile, Celik, and Tian (2021), which is conservative because the strength of the model’s mechanisms is increasing in ϵ .

¹⁸Ideally, η would be the rate at which a firm stops innovating rather than exits, but its effects are small and mostly on the transition path.

Table 1: Externally Calibrated Parameters

Parameter	Description	Value
η	Exit Rate	0.04
γ	Demand Elasticity	3.1
ϵ	Inverse Marginal Innovation Cost Elasticity	1

Internally Calibrated Parameters: I describe the data moments and how they map into parameters. In some cases, the mapping is in closed form, whereas in other cases, I can only describe an intuitive mapping. All moments except the Akcigit and Kerr (2018) patent moment are averages from 1983-1993. This aligns with the first time period in Garcia-Macia, Hsieh, and Klenow (2019), who document the growth rate dynamics I try to explain, as well as the importance of creative destruction for growth and small firm depreciation.

I match the BLS measure of aggregate productivity growth from Garcia-Macia, Hsieh, and Klenow (2019), the sales-weighted average industry revenue share of the largest firm in 4-digit industries in Compustat from Olmstead-Rumsey (2022), and Business Enterprise Expenditures on R&D over GDP from the OECD MSTI database. These are necessary targets because I want to match or explain the changes in these variables over time. Moreover, the large firm’s industry revenue share is an important driver of the model’s mechanisms. I use Compustat rather than Census data because the latter are only available continuously from 1997 onwards, and include many small firms that do not innovate.

The corresponding moments in the model are the growth rate in productivity, the large firm industry revenue share, and aggregate innovation costs over aggregate output. All parameters affect these moments, but the first two moments mostly determine the entry cost, ξ , and the large firm innovation cost shifter, β , which scale the total small firm and large firm innovation costs, respectively. Innovation over output mostly determines firm profits relative to sales given the convexity of the innovation cost function, which pins down the fixed cost $f_S = f_L$ given the markup.

Next, I match the imitation discount, σ , which is the markup, to the cost-weighted average markup estimated in Compustat data by De Loecker, Eeckhout, and Unger (2020). To determine the time discount rate ρ , I match $\rho + g$ to the 1-year real interest rate from

FRED, where g is the growth rate calibrated above.¹⁹

Table 2: Internally Calibrated Parameters and Data Moments

Parameter	Value	Moment	Value
ξ – Entry Cost	38	TFP Growth Rate	1.66%
β – Large Firm Innovation Cost	179.5	Large Firm Revenue Share	40.7%
f_S, f_L – Fixed Cost	0.18	R&D/GDP	1.81%
σ – Imitation Discount	1.3	Markup	1.3
ρ – Time Discount Rate	0.0194	Real Interest Rate	3.6%
λ – Innovation Step Size	1.067	External Innovation Step	0.069
α – Creative Destruction Cost	2.052	Large Job Destruction Rate	25.57%

Finally, two parameters related to creative destruction remain: the innovation step size, λ , and the relative cost, α . To calibrate λ , I match the average growth contribution of a creative destruction innovation, $\frac{\lambda\gamma^{-1}-1}{\gamma-1}$, to the Akcigit and Kerr (2018) estimate of the average growth contribution of an external patent (that mostly cites other firms’ patents).²⁰ The innovation step size is important for the relationship between industry concentration and growth over small firm depreciation. Specifically, a higher λ flattens the relationship between a large firm’s relative innovation rates and its industry revenue share, and between its relative innovation rates and the growth it generates relative to the creative destruction risk it imposes on small firms.

To calibrate the relative creative destruction cost, α , I match an estimate of the rate at which a small firm’s good is creatively destroyed. Given the calibrated growth rate and demand elasticity γ , this determines the creative destruction risk share of small firm depreciation. This share maps the growth a firm generates relative to the *creative destruction risk* it imposes on small firms into growth relative to the *depreciation rate* it imposes on small firms. The latter is what determines the effects of shocks on growth.

To estimate the rate at which a small firm’s good is creatively destroyed, I use the large job destruction rate—the share of aggregate employment lost over a 5 year period

¹⁹This is not the interest rate r given the normalization in Section 4.1.3.

²⁰Their analogous estimate for internal innovations is lower at 0.051, whereas Garcia-Macia, Hsieh, and Klenow (2019) estimate a slightly higher 0.081 using labor flows data through the lens of their growth model.

at firms whose employment shrank by at least two-thirds—computed in Census data by Garcia-Macia, Hsieh, and Klenow (2019). Intuitively, creative destruction generates large job destruction flows because it takes an entire good from a firm instantaneously, whereas in this calibration, growth only depreciates a good’s employment by 16% over 5 years. I compute a *minimum* estimate of the rate at which a small firm good is creatively destroyed, $\bar{\kappa}$, by setting the large job destruction rate equal to $(1 - \exp(-5\bar{\kappa}))L_S^P/(L_S^P + L_S^I)$, which is the gross labor small firms lose to creative destruction as a share of their total labor used for production (L_S^P) and innovation (L_S^I). This produces a minimum estimate for two reasons. First, it supposes all creative destruction of small firm goods in the model generates large negative flows, and is equivalent to supposing a small firm never produces more than one good in a 5 year period. If a small firm has multiple goods, then one can be creatively destroyed without a large negative flow, so more creative destruction is required to hit the target.²¹ Second, it only requires the large job destruction rate *for small firms* in the model to match the data. In the model, large firms do not experience large flows, so the aggregate large job destruction rate is lower than that in the data. Finally, using a minimum estimate is conservative in that it minimizes the strength of the model’s mechanisms.

5 Quantitative Results: A Rise in Concentration

I ask whether a rise in industry concentration driven by an increase in large firm innovation incentives can explain changes in US data since the mid-1990s. The economy starts on the calibrated initial balanced growth path. There is an unanticipated permanent fall in the large firm per-good fixed cost, f_L , to 0.1646 in all industries, calibrated so the new balanced growth path large firm revenue share is 0.51, the 2018 (the latest date available) sales-weighted average revenue share of the largest firm across 4-digit industries in Compustat from Olmstead-Rumsey (2022).²² I interpret this shock as a shift from per-good to firm wide fixed costs—a rise in span of control—due to the rise in information technology.²³

²¹The model does not identify how many goods each small firm produces. I cannot prove this is the minimum across all possibilities because if a firm has multiple goods, then growth can contribute to large negative flows. But that is unlikely to be enough to compensate for the wasted creative destruction.

²²The sales-weighted average R&D share is more volatile, but increases by a similar magnitude from 0.43 in 1993 to a peak of 0.50 in 2011.

²³See Aghion, Bergeaud, Boppart, Klenow, and Li (2022) for a discussion.

I begin by comparing to the data the model results concerning aggregate growth and growth due to different types of innovation. I then inspect the mechanism and contrast the results to what we would get using a similar calibration in the industry model without macroeconomic general equilibrium effects. Next, I show that the model explains a number of related changes in the US economy over a similar time period. Finally, I discuss welfare.

5.1 Industry Concentration and Aggregate Growth

I compare the effects on growth to data from Garcia-Macia, Hsieh, and Klenow (2019) in Table 3. The data are in three consecutive 10-year time periods, which I categorize as before the shock, immediately after the shock, and the new long-run. The fall in the growth rate across balanced growth paths in the model is 35% of the long-run decrease in the data. The rise in the growth rate in the first year after the shock in the model is 120% or 22% of the short-run increase in the data, if I measure model productivity as output over labor, Y_t/\bar{L} , or as Z_t , respectively. The difference is that the former includes a temporary shift in labor from entry to production. Finally, the short-run increase in the data is an average over 10 years, whereas growth in the model is below its pre-shock value after only 4 years.

Table 3: Growth After a Fall in f_L

Moment	Data	Model
Growth Burst	+0.64 ppt (39%) (1993-2003)	GDP: +0.77 ppt (46%) (first year) TFP: +0.14 ppt (9%) (first year)
Cumulative Burst	+6.4 ppt (1993-2003)	GDP: +1.01 ppt (2.5 years) TFP: +0.31 ppt (4 years)
Growth Fall	-0.34 ppt (-20%) (2003-2013)	-0.12 ppt (-7%) (New BGP)

ppt is percentage point rise, and in parentheses is the percent rise. Growth burst is the peak change in the growth rate. Cumulative burst is the peak difference in accumulated growth. Growth fall is the long-run change. GDP and TFP use Y_t/\bar{L} and Z_t to measure productivity, respectively.

Garcia-Macia, Hsieh, and Klenow (2019) also estimate growth in each time period due to firms creatively destroying competitors' goods, and to internal innovation—firms improving

their own goods or developing new ones. I compare the model results to their findings in Table 4, using Z_t for productivity. The model matches that growth from each type of innovation rose and then fell, and that changes in internal innovation were bigger in levels, but smaller as a percentage. The model underestimates initial growth due to creative destruction because I match a minimum estimate of the rate at which a small firm good is creatively destroyed.

Table 4: Growth By Type of Innovation

Moment	Creative Destruction		Internal Innovation	
	Data	Model	Data	Model
Initial	0.44 ppt	0.34 ppt	1.22 ppt	1.32 ppt
Short-Run	+0.24 ppt (55%)	+0.03 ppt (9%)	+0.44 ppt (36%)	+0.11 ppt (8%)
Long-Run	-0.15 ppt (-34%)	-0.03 ppt (-9%)	-0.18 ppt (-15%)	-0.10 ppt (-8%)

Percentage points of growth (in Z_t) and percentage of initial in parentheses, from creative destruction of competitors' goods and internal innovation. Data is from Garcia-Macia, Hsieh, and Klenow (2019). Initial: pre-shock value; 1983-1993. Short-Run: pre-shock value to first year after the shock; 1983-1993 to 1993-2003. Long-Run: across balanced growth paths; 1983-1993 to 2003-2013.

Transition paths for the main outcomes are in Figure 2. The large firm revenue share converges over a similar time interval as the gap between the initial calibration years, 1983-1993, and the target year for the shock, 2018. The growth rate is near its long-run value after 10 years, in line with the empirical time period classification.

5.2 Understanding the Fall in Growth

It is useful to compare the results to what we would find in the industry model from a shock calibrated to match the same rise in concentration. With $\gamma = 1$ and so only creative destruction of competitors' goods, the industry is past the peak of the inverted-U, and growth falls from 1.66% to 1.64%.²⁴ With $\gamma > 1$ and so all types of innovation, the industry is before the peak of the inverted-U, and growth rises from 1.66% to 1.69%.²⁵ In each case, the last

²⁴The initial calibration is the same, except the exogenous measure of small firms replaces the entry cost, and λ is chosen to match large job destruction flows because α is normalized to 1.

²⁵The initial calibration is the same, except the exogenous measure of small firms replaces the entry cost.

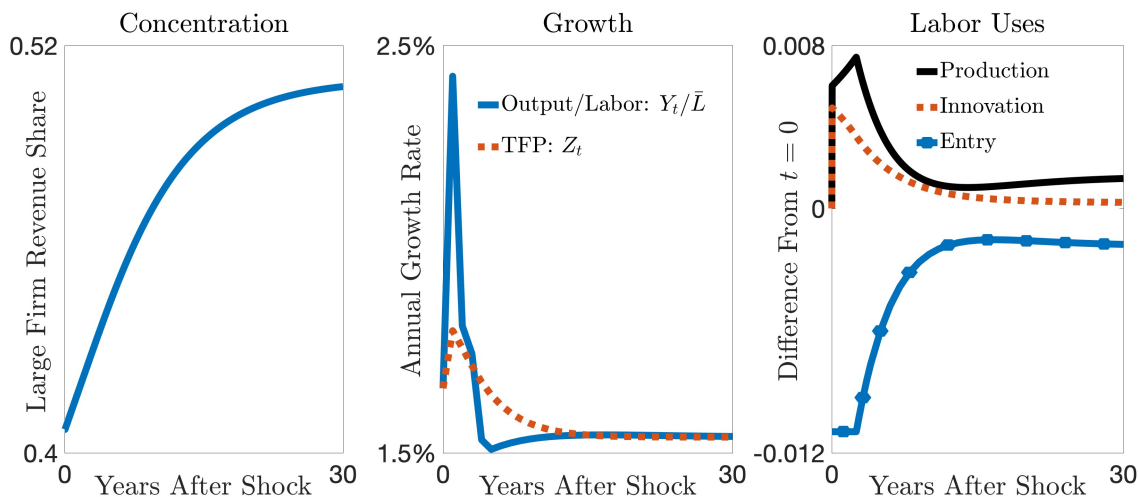


Figure 2: Transition paths after the shock to f_L . Left panel: \mathcal{L}_t . Middle panel: annual growth rates of Y_t/\bar{L} and Z_t . Right panel: labor used for production, innovation, and entry.

panel in Figure 1 displays the long-run relationship between the growth rate and the large firm's industry revenue share following a shock to the large firm's innovation incentives.

In the macroeconomic model, as in the industry model, the growth rate of productivity Z_t is the product of the small firm depreciation rate—the sum of creative destruction risk and relative productivity erosion from growth—and growth relative to small firm depreciation. In the long-run, growth relative to small firm depreciation is the same decreasing function of industry concentration from the industry model with $\gamma > 1$. On the other hand, two macroeconomic forces affect the equilibrium small firm depreciation rate. First, labor shifts from entry to production, which pushes up expenditures on goods, and so the flow profits from producing a good. Second, the free entry condition and a strictly positive entry rate imply that the rate at which each small firm innovates is constant across balanced growth paths because the value of entering is increasing in the benefit of innovating. It follows that in the long-run, the small firm depreciation rate rises to offset exactly the increase in flow profits, and so keep fixed the value a small firm gets from an innovation. Thus, the convex innovation cost effect from the industry model is absent. The result is that following the shock, small firm depreciation and so growth are much lower in the macroeconomic model.

I decompose the time path of the growth rate into small firm depreciation and growth relative to small firm depreciation in the first panel of Figure 3. Growth relative to small firm depreciation falls over time, tracking the rise in concentration. The small firm depreciation

rate rises by more in the short-run than in the long-run because entry hits its lower bound of zero, and the shift in labor from entry to production is bigger in the short-run.

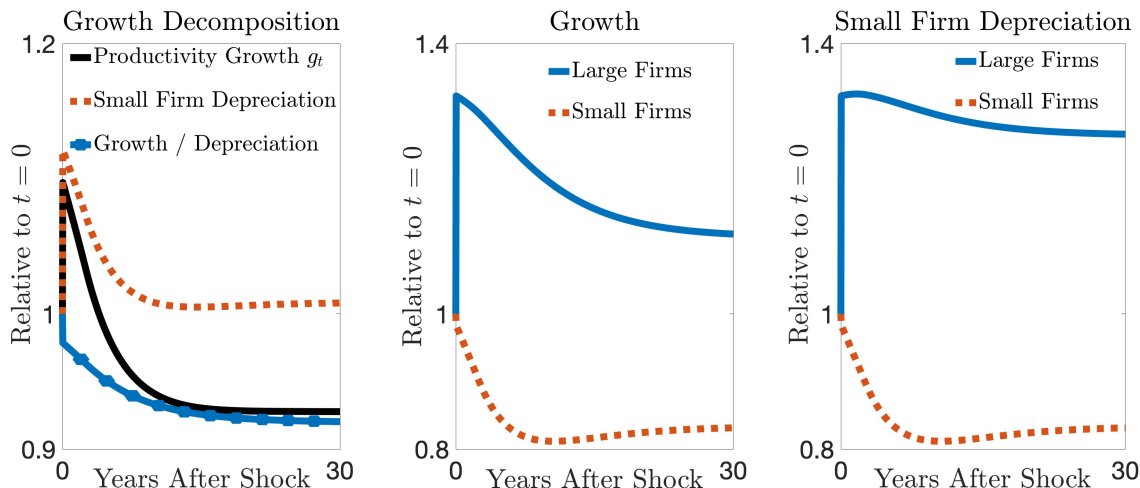


Figure 3: Variables on the transition path relative to their initial values. Left panel: growth in Z_t , the small firm depreciation rate, and growth over small firm depreciation. Middle panel: growth due to large and small firms. Right panel: small firm depreciation rate due to large and small firms.

The second and third panels of Figure 3 display growth and small firm depreciation due to large and small firm innovation. Large firm growth rises immediately after the shock, but returns partly toward its initial value over time because as large firms gain revenue share, they innovate less due to the “Composition Effect” and the “Relative Innovation Effect”. Nonetheless, large firms maintain pressure on small firms because they shift their innovation toward creatively destroying small firm goods. Small firm innovation falls due to the increase in large firm innovation, and falls by more over time as labor shifts partly back from production to entry, and as the measure of small firms decreases.

5.3 Other Features of the Rise in Concentration

R&D Efficiency: In US data, R&D/GDP rose while growth fell. Specifically, R&D/GDP increased from 1.81% in 1983-1993 to 1.91% in 2003-2013, resulting in a fall in the productivity growth rate relative to R&D/GDP from 0.91 to 0.69.²⁶ In the model, the shock explains 34% of this drop in R&D efficiency: across balanced growth paths, R&D/GDP rises from

²⁶Using the same measures of R&D/GDP and growth from the initial calibration.

1.81% to 1.83%, and the growth rate relative to R&D/GDP falls from 0.91 to 0.84. Including entry costs in R&D, growth relative to R&D/GDP falls from 0.63 to 0.60. After the shock, large firms innovate more, and convex costs imply a fall in innovation efficiency.

This result relies on large firms' preference for creatively destroying their competitors' goods. Internal innovation efficiency rises after the shock because large firms are more efficient at internally innovating: their growth generated relative to costs compared to that of small firms is 1.52 and 1.36 in the initial and new balanced growth paths, respectively. The opposite holds for creative destruction: the same ratios are 1.00 and then 0.76. On the one hand, large firm innovation is efficient because they discount the value of generating growth, and so innovate relatively slowly. On the other hand, large firm innovation is inefficient because they have lower fixed costs (after the shock) and face a lower rate of creative destruction, and so innovate relatively quickly. For internal innovation, the first effect dominates even after the shock. For creative destruction, the second effect dominates even before the shock.

Industry Concentration and Industry Growth Rates: Ganapati (2021) finds that across industries in the US, rising concentration is associated with *faster* growth. Specifically, with sector and time fixed effects, a 1 percent rise in the revenue share of the largest 4 firms in a 6-digit NAICS industry over 5 years is associated with 0.1 and 0.2 percent rises in real output and in real output relative to employment, respectively. The model after the shock explains 64% and 36% of these coefficients. Specifically, I run the same regression during the transition path by creating a measure 0 control group of industries with a constant large firm per-good fixed cost. A 1 percent rise in an industry's large firm revenue share over 5 years is associated with 0.06 and 0.07 percent rises in real output and in real output relative to employment (production and innovation), respectively. Intuitively, industry-specific shocks generate a positive relationship between concentration and growth as in the industry model because they do not affect expenditures on goods or the measure of small firms.

Entry: Decker, Haltiwanger, Jarmin, and Miranda (2016) find that in the US, the employment share of firms less than 5 years old fell by 20% during the early 1990s, and then 5% in the early 2000s. In the model, the same share (including production and innovation labor) falls by 17% across balanced growth paths.

5.4 Welfare

Including the transition path, welfare—the present discounted value of consumption—falls by the equivalent of a permanent 3.8% decrease in consumption. Changes in the time paths of growth, labor used for production, and fixed costs contribute -5%, +0.2%, and +1% to this measure, respectively. To see the importance of short-run vs. long-run effects, I compute the first time T so that the discounted value of consumption between 0 and T is lower after the shock than it would have been without the shock. I find that T is 43.24 years.

6 Policy in the Calibrated Model

I use the calibrated model to study changes in the tax rates on large firm revenues, τ_R , and on large firm acquisitions of their competitors' goods, τ_A , as in Sections 3.1.2 and 3.1.3.

6.1 Size-Dependent Taxes

An increase in the tax on large firm revenue, τ_R , that undoes the recent rise in concentration improves welfare. It is strictly better than reversing the fall in the large firm per-good fixed cost, f_L —which would increase welfare as just discussed—because it keeps some of the shift in output from fixed costs to consumption.

In oligopolistic models, it is often optimal to subsidize large firms because they are relatively efficient. Here, innovation efficiency affects welfare directly because it determines how much labor remains for production, and indirectly because if more labor is used for production, then expenditures on goods are higher, which pushes up incentives to innovate. But a new size-related distortion is quantitatively more significant: large firm innovation is tilted toward creatively destroying small firms' goods, which leads to too little innovation. To demonstrate the greater significance of this second effect, I rerun the experiment from Section 5 under the most optimistic scenario for improving the efficiency of innovation: I suppose the increase in labor used for large firm innovation relative to the initial balanced growth path does not subtract from the labor available for other uses.

Welfare and long-run growth still fall, but less than in the original experiment by 45% and 25%, respectively. As expected, there are substantial gains in the efficiency of innovation: growth relative to entry and innovation costs (not including new large firm innovation) rises

by 7% in the long-run, whereas before it fell by 2.9%. This leaves more labor for output, and increases small firm innovation incentives: long-run small firm depreciation rises by 2.2% instead of 0.9%. However, this is not nearly enough to counteract the 7.5% fall in growth relative to small firm depreciation (slightly less than the 8.1% fall before because now the large firm revenue share only rises to 0.503). Intuitively, R&D and entry are small shares of overall labor, so the benefits of improving their efficiency are limited.

6.2 Acquisitions

I study the welfare implications of permanent changes in the tax rate on large firm acquisitions of their competitors' goods, τ_A (see Section 2.2.2 for the model description). I impose that the acquisition search function is isoelastic like the innovation cost functions:

$$\text{for all } x \geq 0, \quad X_A(x) = \omega x^{1/\epsilon_A+1},$$

where $\omega > 0$ is a cost shifter and $\epsilon_A > 0$ (as in Section 3.1.3) is the elasticity of the rate at which a large firm acquires each competitor's good, $A_t(\mathcal{L}_{n,t})$, with respect to the surplus of an acquisition per unit of relative productivity, $V'_t(\mathcal{L}_{n,t}) - (1 + \tau_A)\Pi_t(\mathcal{L}_{n,t})$. I restrict $\tau_A > -1$ because otherwise the large firm offers an infinite price just for the subsidy.

I set the cost shifter sufficiently high so the acquisition rate is small ($A_t(\mathcal{L}_{n,t}) < 0.0001$). This kills non-linear effects of changes in the acquisition rate away from 0, which facilitates comparing results across different specifications. It also makes transition paths irrelevant.

I consistently find that welfare is a U-shaped function of the acquisition tax rate (see the third panel of Figure 4). The intuition follows from the two effects described in Section 3.1.3. First, a lower tax rate increases the acquisition rate, which increases revenue concentration and thus reduces growth. On the other hand, a lower tax rate encourages large firms to innovate less in order to preserve valuable acquisition opportunities, which decreases innovation concentration and thus increases growth. If the tax is just low enough for acquisitions to occur, then the second effect is small because acquisitions are not valuable. At lower values of the tax, the surplus is higher, so the second effect is stronger.

The U-shape of welfare implies the existence of two acquisition tax rate thresholds. First, welfare is higher with acquisitions than without if and only if the tax rate is sufficiently low. Second, the marginal welfare effect of decreasing the tax rate is positive if and only if the tax rate is sufficiently low. I plot these thresholds in the first and second panels of Figure

4, respectively, as functions of the acquisition rate elasticity, ϵ_A . I use the calibration from before the rise in concentration (solid line), after the rise in concentration in Section 5 (dotted line), and after the rise in concentration due instead to a change in the large firm innovation cost shifter β (textured line). The last one eliminates the mechanical increase in the acquisition surplus due to the fall in the large firm per-good fixed cost.

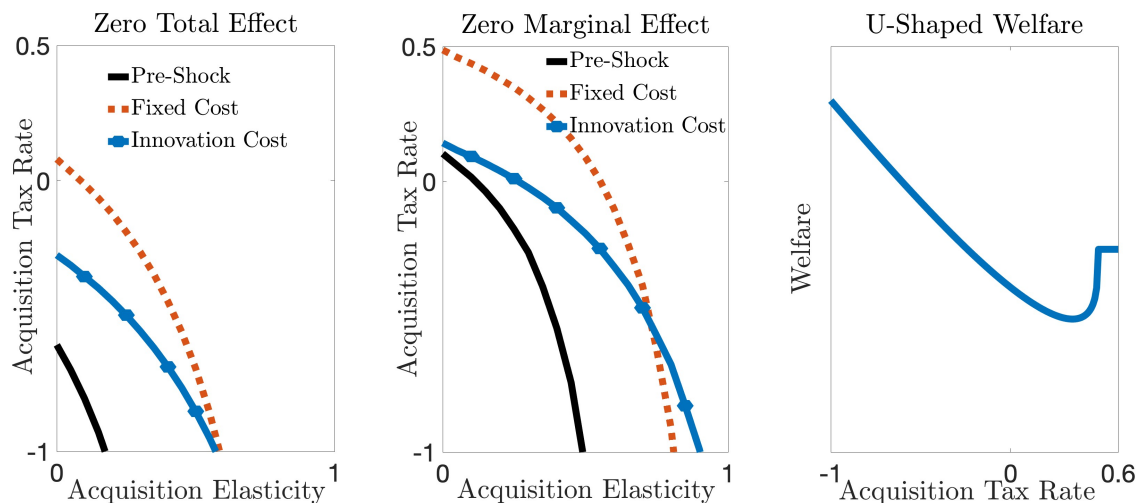


Figure 4: Left panel: τ_A below which welfare improves by allowing acquisitions. Middle panel: τ_A below which welfare improves from a marginal decrease in τ_A . Both thresholds are functions of ϵ_A , and are evaluated using the initial calibration, after the fall in f_L , and after a fall in β calibrated to generate the same rise in concentration. Right panel: example of welfare as a function of τ_A .

The thresholds are falling in the acquisition rate elasticity, ϵ_A , so it is better to raise the acquisition tax if the acquisition rate is more responsive, in line with Theorem 2. Indeed, it is always optimal to raise the tax if the acquisition rate elasticity is higher than the large firm innovation elasticity, $\epsilon = 1$.

After the rise in concentration, the thresholds are higher, so it is better to cut the acquisition tax, in line with Corollary 2. This holds with or without a fall in the large firm per-good fixed cost. As discussed in Section 3.1.3, if concentration is high, then changes in the acquisition rate have little direct effect on growth, but reducing large firm innovation has a big effect because it encourages small firms innovation.

Finally, the second threshold (the marginal effect) indicates whether it is optimal to lower the tax on the margin based on the model mechanisms, allowing for the possibility that the policymaker has motivations outside the model for setting the acquisition tax. The second

threshold is higher than the first, so it may be optimal to lower the tax on the margin even if it is better to ban acquisitions altogether.

7 The Value of Size

When firms can choose the direction of innovation, it is advantageous to be large. Specifically, a large firm gets more value than a small firm from producing a good or from an entrepreneur's innovation technology. We saw this for goods in the previous section because acquisitions occur even with a positive tax rate. I now show that it also applies to innovation technologies. Thus, the theory provides an explanation for mergers and acquisitions absent explicit synergies. Equivalently, it explains why large firms do not break up.

I suppose a large firm has a one-time opportunity to acquire either the internal innovation or creative destruction technology of one of its competitors. I compute the large firm value relative to the small firm value, first holding innovation rates fixed, and then allowing for re-optimization. I use the initial calibration before the rise in concentration so that large and small firms differ only in size. The acquired technologies still exit at rate η .

First, for either innovation technology, holding innovation rates fixed, the large firm is willing to pay 110% of the small firm's value. The large firm values goods more because it faces a lower creative destruction rate. Next, if the large firm can choose innovation rates, it is willing to pay 75% and 179% of the small firm's value for the internal innovation and creative destruction technologies, respectively. In the former case, the value of choosing innovation rates is negative (75% vs. 110%). The large firm cuts the acquired technology's innovation rate by 35% to avoid cannibalization, and anticipating this, small firms innovate more. In the latter case, the value of choosing innovation rates is high (179% vs. 110%). The large firm cuts the acquired technology's creative destruction of large firm goods by 92%, but increases creative destruction of small firm goods by 5%. Anticipating this, small firms innovate less. Finally, the combined effect is that if the large firm can choose innovation rates, it is willing to pay 148% of the small firm's value for both innovation technologies.

These results demonstrate that being large is valuable, but only if firms can choose the direction of their innovation. With internal innovation, there is no direction, and innovation technology is more valuable if held by a small firm. On the other hand, creative destruction innovation technology is more valuable if held by a large firm because the large firm chooses

a different direction as well as magnitude from a small firm: it protects its own goods, but increases pressure on its competitors' goods. This direction choice is valuable because a large firm has many goods to protect.

8 Model Discussion

Constant Markups: Without fixed costs and imitation, the markup on a good would be the minimum of the producer's desired unconstrained markup, and its version's productivity advantage over the second-best version of the good. Given the low calibrated value of γ , the implication would be that when a firm improves a good it already produces, it increases its markup rather than its sales. This would not change the key model result that for a large firm, improving one of its own goods results in more cannibalization than creatively destroying a competitor's good; improving one of its own goods would still increase the ability of other firms to innovate, and would eventually lead to more sales when the good was later creatively destroyed by a competitor. Nonetheless, large firms would be more inclined to improve their own goods because they particularly benefit from raising the average markup.

In support of the model, Argente, Lee, and Moreira (2021) find in US retail scanner data that a good's sales increase and then decrease over time, but its markup is nearly constant.

Finally, it is difficult to analyze the model without constant markups. With constant markups, there is a single industry state, and large firm relative innovation rates are uniquely pinned down. Otherwise, large firm relative innovation rates, or the relative rates at which a large firm creatively destroys different goods with different markups, may for example, be non-transparent ways for large firms to engage in tit-for-tat strategies that deter competition.

Fixed Innovation Technologies: In many endogenous growth models, a firm's ability to innovate scales with its number of past innovations. The empirical findings in Akcigit and Kerr (2018) suggest that this is particularly the case for a firm's ability to internally innovate. I can easily extend the model to include the following innovation technology: if a firm has the most productive version of a good, then it has a unique ability to improve on its version. This would strengthen the model's results by adding a new mechanism to the "Composition Effect" and "Relative Innovation Effect" with the same implications: as revenue shifts to large firms, innovation technology shifts to large firms, which lowers innovation because large firms

internally innovate relatively slowly (as we saw in Section 5.3 on R&D efficiency). Moreover, the values of entering and innovating as a small firm would still depend equally on the small firm depreciation rate, which would be the same function of growth and creative destruction.

Only Small Firms Enter: All entrants are permanently small, so large firm profits do not affect entry. This is important because large firm profits rise in the quantitative experiment in Section 5. One justification for not allowing large firm profits to affect entry is that if large firms exit more slowly or if firms take time to become large, then due to discounting, large firms are over represented in the cross section relative to their salience for a potential entrant. Also, this is the correct approach for the experiment if the interpretation of the shock is that large firms pay higher firm-level costs to achieve lower per-good fixed costs.

9 Conclusion

I develop a model to study how large firms innovate, and the implications for growth and policy. The novel feature of the analysis is that firms choose how much to innovate, as well how much to direct their innovation toward creatively destroying their competitors' goods. I find that a large firm's innovation generates little growth relative to the creative destruction risk it imposes on other firms. This mechanism implies that a rise in large firm innovation incentives can lead to lower overall growth. I characterize conditions so that this is the case. I demonstrate that a rise in large firms' span of control can explain recent changes in the US. Finally, I illustrate the novel implications for size-dependent policies and acquisitions.

Although I study a model in which the interesting decisions concern innovation, the insights likely apply more generally. For example, suppose in a static model of oligopolistic competition, a firm can produce different types of goods, with varying degrees of substitutability with its competitors' goods. The same force that leads larger firms to set higher markups implies that they have a relative preference for producing goods that are close substitutes with their competitors'. Thus, subsidizing large firms to produce more may be suboptimal, unlike in models in which all goods within an industry are equally substitutable.

Bibliography

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. “Innovation, Reallocation, and Growth.” *American Economic Review*, 108(11): 3450-3491.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li. 2022. “A Theory of Falling Growth and Rising Rents.” Working paper.

Aghion, Philippe, Nick Bloom, Richard Blundell, Rachel Griffith, and Peter Howitt. 2005. “Competition and Innovation: An Inverted-U Relationship.” *The Quarterly Journal of Economics*, 120(2): 701-728.

Aghion, Philippe and Peter Howitt. 1992. “A Model of Growth Through Creative Destruction.” *Econometrica*, 60(2): 323-351.

Akcigit, Ufuk, Harun Alp, and Michael Peters. 2021. “Lack of Selection and Limits to Delegation: Firm Dynamics in Developing Countries.” *American Economic Review*, 111(1): 231-275.

Akcigit, Ufuk and William R. Kerr. 2018. “Growth through Heterogeneous Innovations.” *Journal of Political Economy*, 126(4): 1374-1443.

Arrow, Kenneth. 1962. “Economic Welfare and the Allocation of Resources to Invention.” In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, edited by the Universities-National Bureau Committee for Economic Research and the Committee on Economic Growth of the Social Science Research Councils, 609-626. Princeton, NJ: Princeton University Press.

Argente, David, Salomé Baslandze, Douglas Hanley, and Sara Moreira. 2021. “Patents to Products: Product Innovation and Firm Dynamics.” Working paper.

Argente, David, Munseob Lee, and Sara Moreira. 2021. “The Life Cycle of Products: Evidence and Implications.” Working paper.

Atkeson, Andrew and Ariel Burstein. 2019. “Aggregate Implications of Innovation Policy.” *Journal of Political Economy*, 127(6): 2625-2683.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2022. “Labor Market Power.” *American Economic Review*, 112(4): 1147-1193.

- Boar, Corina and Virgiliu Midrigan.** 2022. “Markups and Inequality.” Working paper.
- Broda, Christian and David E. Weinstein.** 2006. “Globalization and the Gains from Variety.” *The Quarterly Journal of Economics*, 121(2): 541-585.
- Cavenaile, Laurent, Murat Alp Celik, and Xu Tian.** 2021. “Are Markups Too High? Competition, Strategic Innovation, and Industry Dynamics.” Working paper.
- De Loecker, Jan, Jan Eeckhout, and Gabriel Unger.** 2020. “The Rise of Market Power and the Macroeconomic Implications.” *The Quarterly Journal of Economics*, 135(2): 561-644.
- De Ridder, Maarten.** 2021. “Market Power and Innovation in the Intangible Economy.” Working paper.
- Decker, Ryan A., John Haltiwanger, Ron S. Jarmin, and Javier 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.
- Edmond, Chris, Virgiliu Midrigan, and Daniel Yi Xu.** 2022. “How Costly Are Markups?” Working paper.
- Fons-Rosen, Christian, Pau Roldan-Blanco, and Tom Schmitz.** 2022. “The Aggregate Effects of Acquisitions on Innovation and Economic Growth.” Working paper.
- Ganapati, Sharat.** 2021. “Growing Oligopolies, Prices, Output, and Productivity.” *American Economic Journal: Microeconomics*, 13(3): 309-327.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow.** 2019. “How Destruction is Innovation?” *Econometrica*, 87(5): 1507-1541.
- Grossman, Gene M. and Elhanan Helpman.** 1991(a). “Innovation and Growth in the Global Economy.” *MIT Press*.
- Grossman, Gene M. and Elhanan Helpman.** 1991(b). “Quality Ladders in the Theory of Growth.” *Review of Economic Studies*, 58(1): 43-61.
- Klette, Tor Jakob and Samuel Kortum.** 2004. “Innovating Firms and Aggregate Innovation.” *Journal of Political Economy*, 112(5): 986-1018.
- Letina, Igor, Armin Schmutzler, and Regina Seibel.** 2021. “Killer Acquisitions and Beyond: Policy Effects on Innovation Strategies.” Working paper.

Olmstead-Rumsey, Jane. 2022. “Market Concentration and the Productivity Slowdown.” Working paper.

Pellegrino, Bruno. 2021. “Product Differentiation and Oligopoly: a Network Approach.” Working paper.

Peters, Michael. 2020. “Heterogeneous Markups, Growth, and Endogenous Misallocation.” *Econometrica*, 88(5): 2037-2073.

Rasmusen, Eric. 1988. “Entry for Buyout.” *The Journal of Industrial Economics*, 36(3): 281-299.

Romer, Paul M. 1990. “Endogenous Technological Change.” *Journal of Political Economy*, 98(5): S71-S102.

Shapiro, Carl. 2012. “Competition and Innovation: Did Arrow Hit the Bull’s Eye?” In *The Rate and Direction of Inventive Activity Revisited*, edited by Josh Lerner and Scott Stern, 361-404. Chicago, IL: University of Chicago Press.

Weiss, Joshua. 2020. “Intangible Investment and Market Concentration.” Working paper.

A Static Equilibrium (Section 2.1.3)

The equilibrium strategies are for each good: 1) only the firm with the most productive version pays the fixed cost; 2) for any first stage actions, the most productive producer in the second stage sets its price equal to the second-most productive producer’s marginal cost, and other producers set their prices equal to their own marginal costs; 3) if multiple firms set the same price, demand is split evenly among those with the lowest marginal cost. I make the following assumption, which always holds in the calibrated model, and insures that the firm with the most productive version of a good finds it optimal to pay the fixed cost.

Assumption 2. $(1 - \sigma^{-1}) R_t > f_S$ and $(1 - (1 - \sigma^{1-\gamma}) \mathcal{L}_{n,t}) (1 - \sigma^{-1} - \tau_R) R_t > f_L$.

In the second stage, a firm without the lowest marginal cost for a good cannot earn positive profits, so it is optimal for them to set price equal to marginal cost. The firm with the lowest marginal cost would set a markup of at least $\gamma/(\gamma - 1)$ if unconstrained by other producers of the good, so by Assumption 1, they are constrained to set a markup of σ .²⁷

²⁷See Edmond, Midrigan, and Xu (2022) for the optimal markup with oligopoly and nested CES demand.

In the first stage, a firm without the most productive version of a good earns zero second stage profits if it pays the fixed cost, so it is optimal not to do so. If a small firm has the most productive version and pays the fixed cost, then it earns second stage profits $\tilde{z}_{n,t}(j)(1 - \sigma^{-1})R_t$, so it is optimal to do so by Assumption 2. If the large firm has the most productive version, then it takes into account that paying the fixed cost decreases the good's equilibrium price, and so reduces the sales of its other goods. If the large firm pays fixed costs for a share x of its relative productivity, then its industry revenue share is

$$\tilde{\mathcal{L}}_{n,t}(x) \equiv x\mathcal{L}_{n,t}/(1 - (1 - x)(1 - \sigma^{1-\gamma})\mathcal{L}_{n,t}),$$

and its profit across both stages is $\tilde{\mathcal{L}}_{n,t}(x)(1 - \sigma^{-1} - \tau_R)R_t - x\mathcal{L}_{n,t}f_L$, which reflect that if the large firm does not pay the fixed cost for a good, then that good's price rises by a factor of σ . Profit is strictly concave in x , so it is optimal for the large firm to play the equilibrium strategy if the first derivative at $x = 1$ is positive, which is Assumption 2.

B Proofs for Section 3.1.2

Proof of Theorem 1. On a balanced growth path, the growth rate is $g = (\kappa_L + N\kappa_S)(1 - \mathcal{L}^2)$. HJB equation (9) for a small firm's value of a good yields $\Pi = (1 - \sigma^{-1})R/(r + N\kappa_S + \kappa_L)$. Plugging into the small firm innovation first order condition implies

$$WX'_S(\kappa_S) = \frac{1}{\ln(\lambda)} \frac{(1 - \sigma^{-1})R}{r + N\kappa_S + \kappa_L},$$

which yields κ_S , and so \mathcal{L} and g as differentiable functions of equilibrium κ_L . In particular,

$$\frac{\partial N\kappa_S}{\partial \kappa_L} = \frac{-N\kappa_S\epsilon_S}{r + \kappa_L + N\kappa_S(1 + \epsilon_S)} \quad \frac{\partial \mathcal{L}}{\partial \kappa_L} = \frac{1}{\kappa_L + N\kappa_S} \left(1 - \mathcal{L} \frac{r + \kappa_L + N\kappa_S}{r + \kappa_L + N\kappa_S(1 + \epsilon_S)} \right),$$

where $\epsilon_S = X'_S(\kappa_S)/(\kappa_S X''_S(\kappa_S))$ is the elasticity of κ_S with respect to the expected discounted profits a small firm gets from a good. It follows that

$$\frac{\partial g}{\partial \kappa_L} = \frac{r + \kappa_L + N\kappa_S}{r + \kappa_L + N\kappa_S(1 + \epsilon_S)}(1 - \mathcal{L}^2) - 2\mathcal{L} \left(1 - \mathcal{L} \frac{r + \kappa_L + N\kappa_S}{r + \kappa_L + N\kappa_S(1 + \epsilon_S)} \right),$$

where the first term is the positive derivative of the total rate at which a small firm good is creatively destroyed, and the second term is the negative derivative of growth relative to that rate. Some algebra and dividing by $2\mathcal{L}(1 - \mathcal{L})\frac{\kappa_L + N\kappa_S}{r + \kappa_L + N\kappa_S(1 + \epsilon_S)}$ yields the result. ■

Proof of Corollary 1. The corollary follows from the proof of Theorem 1 above, which shows that equilibrium $\kappa_L + N\kappa_S$ and \mathcal{L} are increasing functions of equilibrium κ_L . ■

C Equilibrium Uniqueness and Stability

Proposition 1. *If $\epsilon_L \epsilon_S / (1 + \epsilon_S) < 1$, then there is a unique equilibrium, where ϵ_L is the elasticity of κ_L with respect to the large firm value of a good, V' , and ϵ_S is the elasticity of κ_S with respect to the small firm value of a good, Π . In that case, if there are no acquisitions ($A = 0$), then equilibrium κ_L is a differentiable increasing function of $((1 - \sigma^{-1} - \tau_R)R - f_L) / \beta$.*

Proof. Use HJB equation (7) to solve for the equilibrium large firm value of a good:

$$V' = \frac{(1 - \sigma^{-1} - \tau_R)R - f_L + W \ln(\lambda) \beta X_L(\kappa_L) + W X_A(A) + A(1 + \tau_A)\Pi}{r + \kappa_L + A + N\kappa_S}, \quad (16)$$

where Π is given in the proof of Theorem 1. The large firm first order conditions are

$$W \beta X'_L(\kappa_L) = V' / \ln(\lambda) \quad W X'_A(A) = \max\{V' - (1 + \tau_A)\Pi, 0\},$$

which yields κ_L and A as differentiable functions of equilibrium V' and Π . Moreover, the proof of Theorem 1 showed that small firm optimization yields κ_S and so Π as differential functions of equilibrium κ_L . It follows that these variables are all differentiable functions of equilibrium κ_L . Thus, let $T(V')$ be the differentiable right-hand side of (16). To prove the first part of the proposition, it is sufficient to show there is a unique V' such that $T(V') = V'$.

First, $T(\cdot) \geq (1 - \sigma^{-1} - \tau_R)R - f_L > 0$. It follows that $V' > 0$ and that as V' goes to 0, $T(V') > V'$. Hence, it is sufficient to show that if $T'(\cdot) < 1$. To compute $T'(V')$, note that all effects of changes in κ_L and A sum to 0. Thus,

$$T'(V') = \frac{A(1 + \tau_A) \frac{\partial \Pi}{\partial V'} - V' \frac{\partial \kappa_S}{\partial V'}}{r + \kappa_L + A + N\kappa_S} = \frac{-A(1 + \tau_A)\Pi + V' N \kappa_S \epsilon_S}{(r + \kappa_L + A + N\kappa_S)(r + \kappa_L + N\kappa_S(1 + \epsilon_S))} \frac{\kappa_L}{V'} \epsilon_L,$$

where the second equality follows from the proof of Theorem 1 and the large firm innovation first order condition, and where $\epsilon_I = X'_I(\kappa_I) / (\kappa_I X''_I(\kappa_I))$ for $I \in \{S, L\}$. Since $\tau_A > -1$ and $r > 0$, it follows that $T'(V') < \epsilon_L \epsilon_S / (1 + \epsilon_S)$. The first part of the proposition follows.

For the second part of the proposition, set $A = 0$, divide both sides of (16) by β , divide both sides of the large firm innovation first order condition by β , and define $\bar{V}' \equiv V' / \beta$ to get that in equilibrium, these equations are

$$\bar{V}' = \frac{((1 - \sigma^{-1} - \tau_R)R - f_L) / \beta + W \ln(\lambda) X_L(\kappa_L)}{r + \kappa_L + N\kappa_S} \quad W X'_L(\kappa_L) = \bar{V}' / \ln(\lambda). \quad (17)$$

As in the proof of the first part of the proposition, there is a unique equilibrium \bar{V}' if $\epsilon_L \epsilon_S / (1 + \epsilon_S) < 1$. Moreover, in that case, the derivative of the right-hand side of the first

equation in (17) is strictly less than 1. It follows that if $((1 - \sigma^{-1} - \tau_R)R - f_L)/\beta$ increases, then equilibrium \bar{V}' must increase for the first equation in (17) to continue to hold. Thus, from the second equation in (17), κ_L must increase. ■

The proof of the first part of the proposition shows that if there are no acquisitions ($A = 0$), then a necessary condition for an equilibrium to be unique is

$$\frac{\kappa_L N \kappa_S \epsilon_S}{(r + \kappa_L + N \kappa_S)(r + \kappa_L + N \kappa_S(1 + \epsilon_S))} \epsilon_L \leq 1.$$

Otherwise, the equilibrium is unstable in the sense that if the large firm innovates marginally faster (κ_L increases), then small firms react so much ($\partial \kappa_S / \partial \kappa_L$ is sufficiently negative) that the large firm's marginal benefit of innovating even faster strictly exceeds the marginal cost.

D Proofs for Section 3.1.3

Proof of Theorem 2. An approximately linear small firm innovation cost function, $X_S(\cdot)$, is the limit as the elasticity ϵ_S of small firm innovation with respect to the small firm value of a good goes to infinity. It follows from the proof of Theorem 1 that changes in the acquisition tax rate, τ_A , do not affect the equilibrium total rate at which a small firm good is creatively destroyed, $\kappa_L + N \kappa_S$, or small firm value of a good, Π . It follows from (10) that

$$\frac{\partial g}{\partial \tau_A} = \left(-\mathcal{L} \frac{\partial \kappa_L}{\partial \tau_A} - \frac{\kappa_L}{\kappa_L + A + N \kappa_S} \left(\frac{\partial \kappa_L}{\partial \tau_A} + (1 - \mathcal{L}) \frac{\partial A}{\partial \tau_A} \right) \right) \ln(\lambda), \quad (18)$$

where the derivatives are of equilibrium outcomes with respect to τ_A . From the large firm first order conditions and since $A > 0$, it follows that

$$\frac{\partial \kappa_L}{\partial \tau_A} = \frac{\kappa_L \epsilon_L}{V'} \frac{\partial V'}{\partial \tau_A} \quad \frac{\partial A}{\partial \tau_A} = \frac{A \epsilon_A}{V' - (1 + \tau_A) \Pi} \left(\frac{\partial V'}{\partial \tau_A} - \Pi \right).$$

The proof of the first part of Proposition 1 for a unique equilibrium yields

$$\frac{\partial V'}{\partial \tau_A} = \frac{A \Pi / (r + \kappa_L + A + N \kappa_S)}{1 - T'(V')} = \frac{A \Pi}{r + \kappa_L(1 - \epsilon_L) + A + N \kappa_S},$$

where I take the limit as ϵ_S goes to infinity in the equation for $T'(V')$. Plugging into (18) and multiplying both sides by

$$\frac{\kappa_L + A + N \kappa_S}{\kappa_L(1 - \mathcal{L})A(r + \kappa_L(1 - \epsilon_L) + N \kappa_S)} \frac{V' - (1 + \tau_A) \Pi}{\Pi},$$

which is strictly positive, completes the proof of the theorem. ■

Proof of Corollary 2. The corollary follows from the proofs of Theorems 1 and 2, which show that the large firm value of a good, V' , is strictly decreasing in both f_L and τ_R . Moreover, given $\epsilon_S \rightarrow \infty$, they show that κ_L and A are strictly increasing in V' , and that $\kappa_L + N\kappa_S$ and Π are constant in V' . ■

E Proof of Theorem 3

First, \mathcal{L} is strictly increasing in $\kappa_L(S)$ because otherwise if $\kappa_L(S)$ goes up, then other types of large firm innovation definitely increase as well, and so the large firm's industry revenue share has to rise. Moreover, as $\kappa_L(S) \rightarrow 0$, \mathcal{L} goes to 0, and as $\kappa_L(S) \rightarrow \infty$, \mathcal{L} goes to 1.

Before proving the four items in the theorem, I describe key features of the full model with $\gamma > 1$. The small firm depreciation rate is $\kappa_L(S) + N\kappa_S + (\gamma - 1)g$. Let \tilde{g}_S be the rate at which a small firm generates growth relative to the depreciation it imposes on small firms, and let $\tilde{g}_L(\mathcal{L})$ be the ratio for the large firm. From the large firm relative innovation rates, $\tilde{g}'_L(\mathcal{L}) < 0$, $\tilde{g}_L(0) = \tilde{g}_S$, and $\tilde{g}_L(1) = 0$. The long-run growth rate is

$$g = (\kappa_L(S) + N\kappa_S + (\gamma - 1)g)((1 - \mathcal{L})\tilde{g}_S + \mathcal{L}\tilde{g}_L(\mathcal{L})).$$

Finally, small firm optimization yields

$$(1 - \mathcal{L})(\kappa_L(S) + N\kappa_S + (\gamma - 1)g) = A(r + \kappa_L(S) + N\kappa_S + (\gamma - 1)g)^{-\epsilon}, \quad (19)$$

where $A > 0$ is a constant, the left-hand side is the depreciation rate imposed on a small firm by other small firms, which scales linearly with κ_S , and the right-hand side other than A reflects that the value a small firm gets from a good depends inversely on the small firm discount rate, which is the sum of the interest rate and the small firm depreciation rate.

For the first item in the theorem, as $\kappa_L(S)$ goes to 0, $\partial g / \partial \kappa_L(S)$ goes to the derivative of the small firm depreciation rate with respect to $\kappa_L(S)$ because \mathcal{L} goes to 0 and $\tilde{g}_L(0) = \tilde{g}_S$. Since \mathcal{L} is strictly increasing in $\kappa_L(S)$, it follows from (19) that the small firm depreciation rate is strictly increasing in $\kappa_L(S)$.

For the second and third items, as $\kappa_L(S)$ goes to infinity, the small firm depreciation rate goes to infinity. It follows from (19) that it converges to the product of a strictly positive constant and $(1 - \mathcal{L})^{-1/(\epsilon+1)}$. Using the large firm relative innovation rates, the growth rate relative to the small firm depreciation rate converges to the product of a strictly

positive constant and $(1 - \mathcal{L})^{\min\{1, \epsilon\}}$, where the 1 comes from small firm innovation and large firm creative destruction of small firm goods, and the ϵ comes from other types of large firm innovation. Thus, growth converges to the product of a positive constant and $(1 - \mathcal{L})^{\min\{\frac{\epsilon}{\epsilon+1}, \frac{\epsilon^2 + \epsilon - 1}{\epsilon + 1}\}}$. If $\epsilon > 1$, then the minimum is the first option, which is strictly positive. If $\epsilon < 1$, then the minimum is the second option, which is strictly positive if and only if $\epsilon > (\sqrt{5} - 1)/2$. The result follows.

For the fourth item, set $\epsilon = 1$. Given the first and second items, it is sufficient to show that there $\partial g / \partial \kappa_L(S)$ is decreasing in $\kappa_L(S)$. From (19), the derivative of small firm depreciation with respect to $\kappa_L(S)$ is

$$\frac{(\kappa_L(S) + N\kappa_S + (\gamma - 1)g)(r + \kappa_L(S) + N\kappa_S + (\gamma - 1)g)}{(1 - \mathcal{L})(r + 2(\kappa_L(S) + N\kappa_S + (\gamma - 1)g))} \frac{\partial \mathcal{L}}{\partial \kappa_L(S)}.$$

From large firm relative innovation rates, the derivative of growth relative to small firm depreciation with respect to \mathcal{L} is

$$\left(\tilde{g}_L(\mathcal{L}) - \tilde{g}_S - \frac{\mathcal{L}}{1 - \mathcal{L}} \frac{1}{2\lambda^{\gamma-1}(1 - \mathcal{L}) + \mathcal{L}} \tilde{g}_L(\mathcal{L}) \right) \frac{\partial \mathcal{L}}{\partial \kappa_L(S)},$$

where the last term is $\mathcal{L}\tilde{g}'_L(\mathcal{L})$ using $\tilde{g}_L(\mathcal{L}) = (\gamma - 1)^{-1}\lambda^{\gamma-1}(1 - \mathcal{L})/(2\lambda^{\gamma-1}(1 - \mathcal{L}) + \mathcal{L})$. Thus, multiplying $\partial g / \partial \kappa_L(S)$ by $(1 - \mathcal{L})/g$, dividing by $\partial \mathcal{L} / \partial \kappa_L(S)$, subtracting one from the first term (the product of growth relative to small firm depreciation and the derivative of small firm depreciation), and adding one to the second term (the product of small firm depreciation and the derivative of growth relative to small firm depreciation) yields

$$\frac{-(\kappa_L(S) + N\kappa_S + (\gamma - 1)g)}{r + 2(\kappa_L(S) + N\kappa_S + (\gamma - 1)g)} + \frac{\tilde{g}_L(\mathcal{L})}{(1 - \mathcal{L})\tilde{g}_S + \mathcal{L}\tilde{g}_L(\mathcal{L})} \frac{2\lambda^{\gamma-1}(1 - \mathcal{L})}{2\lambda^{\gamma-1}(1 - \mathcal{L}) + \mathcal{L}}.$$

This expression is decreasing: the first term is decreasing because small firm depreciation is increasing in $\kappa_L(S)$; the first fraction in the second term is decreasing because $\tilde{g}_L(\mathcal{L})$ is decreasing and is less than \tilde{g}_S ; the second fraction in the second term is decreasing because \mathcal{L} is increasing in $\kappa_L(S)$. The result follows.