Market Power when Ideas get Harder to Find: A Theory of Directed Innovation

Julian Schärer *

February 2023

Abstract

I provide a theory that connects the rise of market power (De Loecker, Eeckhout, and Unger 2020) to the decline in research productivity (Bloom et al. 2020). I build a Schumpeterian firm dynamics model with directed innovation and show that the directedness of innovations has macroeconomic implications: Directed innovation leads to selection of firms based on market power. The model predicts rising markups when *ideas are getting harder to find*: As the stakes in the innovation process rise, the incentives to direct innovations increase; this leads to stronger selection on market power and an upward shift of the whole markup distribution. Moreover, despite all moments being untargeted, the model predictions on factor shares, average firm size, productivity growth and R&D-to-GDP ratio are all in line with the data.

1 Introduction

For decades, markups have been on the rise (De Loecker, Eeckhout, and Unger 2020) while ideas have gotten harder to find (Bloom et al. 2020). This is the first paper that connects these two striking macroeconomic trends. To formalize the argument, I develop a theory of directed innovation. The theory predicts that as research productivity declines, innovators direct their innovation efforts away from markets owned by dominant firms – consequently, dominant firm are more likely to survive, thereby increasing the level of market power in the economy.

*Julian Schärer, Department of Economics, University of Zurich. E-mail: julian.schaerer@econ.uzh.ch.
I am grateful to Daron Acemoglu, Dave Donaldson, Alessandro Ferrari, Bernhard Ganglmair, David Hémous, Severin Lenhard, Paolo Mengano, Ralph Ossa, Michael Peters, Florian Scheuer, Armin Schmutzler, and Christian Wolf for valuable discussions and comments.

The key mechanism for this result is directed innovation. I define directed innovation as the process of *searching for* and *innovating in* an attractive market. The directedness of innovation would not matter if creative destruction were frictionless and innovators could always improve upon the existing technological frontier. In reality, however, there are several *barriers to frontier innovation*: Companies have trade secrets, technology diffusion is sluggish because knowledge does not flow freely, and incumbents may employ strategic patents to prevent entry. Dominant firms are less likely to be replaced, giving rise to a notion of dynamic market power.

The relevance of the mechanism is supported by recent evidence indicating that market power is self-reinforcing. Argente et al. (2020) show that patenting by market leaders leads to less innovation by competitors and entrants. In a similar vein, Galasso and Schankerman (2015) find that if patent rights of large firms are removed by court invalidation, there are more patent citations building on the focal patent. This effect begins only two years after the focal patent has been invalidated, which is consistent with the entry of new firms that engage in cumulative innovation (Galasso and Schankerman 2015). In this case, the removal of technological dominance (and thus market power) spurs innovative activity.

A natural interpretation of the above findings is that innovators have incentives to direct their innovation efforts towards markets where they will be competing with a ceteris paribus less dominant leader. As a consequence, dominant firms are better shielded against creative destruction, leading to selection of firms based on market power. This selection effect, which I term *Schumpeterian selection*, determines the economy-wide level of market power. The first goal of this paper is to incorporate directed innovation and Schumpeterian selection effects into an otherwise standard framework of quality-ladder firm dynamics. The key prediction of the model is that the degree of directed innovation determines the degree of Schumpeterian selection, which in turn pins down the distribution of markups and the average economy-wide markup.

I build upon the Schumpeterian firm dynamics framework of Klette and Kortum (2004) and develop a Schumpeterian theory of innovation frictions, directed innovation, and market power. The *barriers to frontier innovation* are summarized by a single parameter: the probability that an innovator cannot improve upon the leading technology (and has to build upon the second-best technology instead).

The novel concept of directed innovation entails market search and an innovation effort directed to the most attractive market in randomly searched sample. Search is

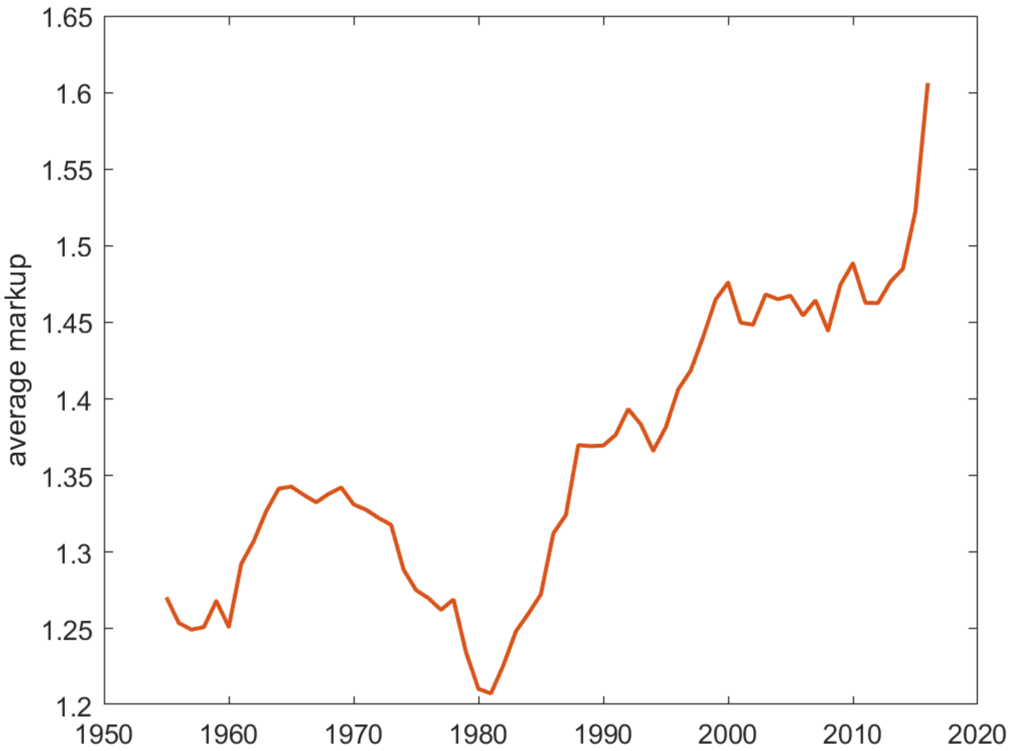

Figure 1: Average markup in the US. Source: Data from De Loecker, Eeckhout, and Unger (2020)

costly, giving rise to a trade-off that balances the cost of search and the cost of R&D. Market search allows innovators (both incumbents and entrants) to sample more potential markets, enabling them to direct their innovation efforts towards more attractive markets.

Innovating in a market of a less dominant (i.e. low markup) leader is desirable: While it is hard to overcome a large leader-follower gap, innovation success is virtually guaranteed in a small leader-follower gap market. Given the implied imbalances in appropriability of innovations, innovators have incentives to direct their efforts towards markets where they face – ceteris paribus – less dominant leaders. This, in turn, provides a rationale for market search.

Since dominant firms that strengthen their leadership with patents are often firms with a lot of market power, high-markup firms are less likely to be targeted by innovators – and thus less likely to be replaced. This gives rise to Schumpeterian selection effects: Dominance in the product market shields from future competition, so firms with above-average markups tend to grow large relative to firms that are exposed to more competitive pressure.

I find that the degree of innovation directedness is indeed a key variable for the macroeconomy: it determines how much selection of firms based on market power there is, which in turn affects the distribution of market power, the level of monopolistic markup distortion, the average markup, and factor shares of the economy.

What are the model predictions in an environment where ideas get harder to find? In section 4, I calibrate the model to data from the US pre-1980 economy; in section 5, I then study the implications of a decline in research productivity. As R&D becomes costlier, the stakes in the innovation process rise. Anticipating the greater cost of R&D, innovators are willing to allocate more resources towards market search in order to better direct their R&D efforts. Compared to an initial steady state, high-markup firms are targeted even less, which shields them even more from creative destruction. The greater degree of Schumpeterian selection implies rising market power over time – as has been observed empirically (see figure 1 and De Loecker, Eeckhout, and Unger 2020). Moreover, the profit share increases at the expense of the labor share, again mirroring the data trend.

It is worth emphasizing that by definition (the decline in research productivity is simply taken from Bloom et al. 2020), all moments are untargeted. Other predictions of the model are also in line with the data: While declining research productivity itself disincentivizes R&D, this is more than compensated by higher expected profits over the firm life cycle due to higher average markups – the latter effect incentivizes R&D sufficiently so that R&D-to-GDP expenditures increase. The rise in R&D-to-GDP is, however, still dominated by the decline in research productivity, meaning that productivity growth slows down. Further, provided that entrant innovation is more elastic than incumbent innovation [1], the number of entrant innovations is falling relatively more than the number of incumbent innovations; as a consequence, the model not only predicts a falling firm entry rate but also a rise in average firm size.

The model makes a number of testable predictions: (i) the persistence of product market leadership should have increased over time; (ii) there is a positive relationship between markups and survival rates; and (iii) this relationship should have gotten stronger over time.

1. Which is, for instance, trivially satisfied if there is no entrant congestion externality; in this case the entrant R&D cost is a constant.

Related literature and main contributions

While the modelling tools introduced by this paper are novel, it is well known that the social value of innovations often exceed their private values because of limited appropriability of innovations (Nelson 1959, Arrow 1962). Yet, for reasons of analytical tractability, the literature on economic growth has mostly abstracted from the interaction between the market-specific state of competition and the directedness of innovation.

In their seminal contributions, Aghion and Howitt (1992) and Grossman and Helpman (1991) have formalized the paradigm of creative destruction that drives economic growth. The Schumpeterian firm dynamics literature – starting with Klette and Kortum (2004) – has subsequently built upon the seminal work from the 1990s. However, innovations are still undirected and take place randomly in one of the markets. On the other hand, the fruitful strand of step-by-step innovation models following the tradition of (Aghion, Harris, and Vickers 1997) has studied both theoretically and empirically how incumbent firms within an industry invest strategically in R&D depending on their relative technological level within the industry. A key feature of the step-by-step framework is that the incentives to innovate are decreasing in the dominance of the leading firm. However, these models chiefly miss important aspects of firm dynamics, as incumbent firms do neither face creative destruction by entrants nor expand into new markets.

The model I develop can account for market-specific innovation incentives while retaining the Schumpeterian entry and exit dynamics. From the incumbents' perspective, the model gives rise to firm dynamics that feature selection based on market power.

In this paper, I make two main contributions: (i) I introduce a notion of directed innovation into the Schumpeterian firm dynamics framework; and (ii) I provide an explanation why we observe rising markups as ideas get harder to find. The model can also account for other trends like the decline in productivity growth, the fall of the labor share Elsby, Hobijn, and Sahin (2013), the fall of the firm entry rate (R. A. Decker et al. 2016), the declining share of activity from young firms (R. Decker et al. 2014) and the increase in average firm size (Hopenhayn, Neira, and Singhania 2018).

The mechanism of directed innovation and Schumpeterian selection is crucial for all model predictions. If this channel were shut off, e.g. in the absence of *barriers to frontier innovation*[2], the framework could not account for any empirical trend other than the falling rate of aggregate productivity growth; evidently, the simplifying assumption

2. In the absence of barriers to frontier innovation, incentives to direct innovations are absent and there is no Schumpeterian selection.

of undirected innovation – albeit common in the literature –, would mask important dynamics.

Other explanations for some of the macroeconomic trends include the role of IT (Aghion et al. 2021); the role of intangibles (De Ridder 2020); declining knowledge diffusion between leaders and followers (Akcigit and Ates 2021); declining innovativeness of small firms (Olmstead-Rumsey 2022); declining interest rates (Liu, Mian, and Sufi 2019); and declining population growth (Peters and Walsh 2021). Manera (2021) builds a Schumpeterian growth model where monopolistic firms conduct defensive patenting, giving rise to misallocation of innovative talent.

The remainder of the paper is organized as follows. Section 2 introduces a Schumpeterian theory of directed innovation and section 3 discusses the implications of a decline in research productivity through the lens of the model. Section 4 presents a calibration of the US pre-1980 economy and section 5 compares the quantitative predictions of the 2010s economy and the pre-1980s economy of the United States. Section 6 concludes.

2 Model

In this section, I first present the model setup and describe the theoretical novelties (barriers to frontier innovation, market search and directed innovation). Then, I characterize the static equilibrium, the balanced growth path, and important comparative statics.

Preferences The representative consumer has logarithmic preferences,

$$U_0 = \int_0^\infty e^{-\rho t} \ln(C_t)\, dt,$$

where C_t is consumption of the final good at time t and ρ is the time discount rate. The price of the consumption good is normalized to one so that the budget constraint of the representative consumer is given by

$$C_t + \dot{A}_t = w_t L_t + r_t A_t,$$

where A_t denotes total assets and $L_t \equiv 1$ is total labor which is constant and supplied inelastically. The wage rate w_t and the interest rate r_t are the relevant relative prices. Normalizing the price of the final good to unity, the Euler equation at any time t is given by

$$r_t = \rho + \frac{\dot{C}_t}{C_t} \tag{1}$$

Further, the following transversality condition holds:

$$\lim_{t\to\infty} \exp\Big(- \int_0^t r_t \, dt \Big) A_t = 0.$$

Final Good The perfectly competitive final good sector transforms differentiated intermediate products into the final good according to

$$Y_t = \exp\Big(\int_0^1 \ln y_{it} \, di \Big), \tag{2}$$

where y_{it} is the amount of intermediate variety i that is sourced from the intermediate firm offering it at the lowest price.

All of the final good is consumed so that

$$C_t = Y_t.$$

holds at each point in time t.

Intermediate Goods Intermediate producers hire labor to produce products. The leader in product line i with labor productivity q_{it} produces according to

$$y_{it} = q_{it} l_{it},$$

where l_{it} denotes the labor input. All other firms have access to the public technology $q_{-it} < q_{it}$, which restricts the leader's pricing. Technological progress is labor-saving and in any product line, the different brands of the leader and the fringe firms are perceived as perfect substitutes.

Pricing In each product line i, the leader and the fringe firms engage in Bertrand competition, and if the leader offers the same price as the fringe firms, the final good sector purchases variety i entirely from the leader. In equilibrium, limit pricing by the leader leads to prices equal to the marginal cost of the fringe firms. From the maximization problem of the representative final good producer, it then follows that all intermediate products are sources from the respective leader and that the same amount is spent on each intermediate variety. Thus, product demands are

$$y_{it} = \frac{Y_t}{p_{it}},$$

where p_{it} is the lowest price offered for product i.

Free Entry There is free entry into market exploration. Entry requires hiring two types of workers – analysts to search potential markets and R&D workers to generate an idea. An entrant chooses how many randomly drawn markets $s \geq 1$ to explore simultaneously, which requires hiring $\frac{1}{\phi}\eta s^{\frac{1}{\eta}}$ analysts at the market-clearing wage. The parameter η denotes the search elasticity, ϕ the search productivity, and s is the search intensity. If search were sufficiently costly, i.e. if ϕ is sufficiently small, the corner solution $\tilde{s} = 1$ would be binding. In this case, there is only minimal search and no directed innovation, thus nesting the canonical quality-ladder model of undirected innovation (Grossman and Helpman 1991 and Klette and Kortum 2004). However, in the more plausible case, search is not prohibitively costly and $s^* > 1$ holds, meaning that the margins of search and directed innovation are relevant. In the calibration in section 4, s^* (and by consequence ϕ) is identified from empirical moments and the estimated value is approximately 2 for the initial steady state.

At a Poisson rate of 1, a potential entrant receives a signal indicating which of the s markets is most attractive for entry – which corresponds to the market with the least dominant current leader charging the *lowest* markup. Markets with currently low markups are more attractive for entrants because in these markets, entrants are more likely to surpass the leaders and establish themselves as incumbents at the technological frontier. While exploring more markets simultaneously yields higher-quality signals, this becomes increasingly costly.

Once an entrant has received a signal in which market to innovate, she hires $\frac{1}{\theta}(\mathcal{M}_{0,t})^{\delta}$ R&D scientists at the market-clearing wage w_t in order to generate an innovation. The R&D cost is weakly increasing in the total number of potential entrants $\mathcal{M}_0$, which can be interpreted as a congestion force (if $\delta > 0$).

In summary, the total entry costs are

$$L_{entry}(s|\mathcal{M}_0) = \underbrace{\frac{1}{\phi}\eta s^{\frac{1}{\eta}}}_{\text{search}} + \underbrace{\frac{1}{\theta}(\mathcal{M}_0)^{\delta}}_{\text{R\&D}}.$$

Incumbent firms Firm dynamics are similar as in Klette and Kortum (2004). Firms expand by innovating in product lines currently operated by other firms, and shrink if another firm innovates in one of their product lines. While new firms enter with a single product, one-product firms that lose their last product exit permanently.

To search a sample of s markets arriving at rate x, a firm needs to employ L_S analysts, where

$$L_S(s, x) = x \frac{1}{\phi} \eta s^{\frac{1}{\eta}}.$$

Thus, the search technology is the same as for entrants, adjusted for the innovation rate x (which is always equal to one in the case of entrants). This property implies that the optimal search intensity of both entrants and incumbents is the same and model remains tractable.

To innovate at rate x, a size-n firm needs to employ $L_{R\&D}$ R&D scientists,

$$L_{R\&D}(x|n) = \frac{1}{\tilde{\theta}} \gamma x^{\frac{1}{\gamma}} n^{\frac{\gamma-1}{\gamma}}$$

The normalization $n^{\frac{\gamma-1}{\gamma}}$ ensures that the optimal expansion rate x is proportional to firm size as in Klette and Kortum (2004). Combining, a firm of size n that searches s markets and innovates at rate x has to hire $L_S(s, x) + L_{R\&D}(x|n)$ units of labor.

Innovation Process Upon receiving a signal of the most attractive market, the innovator – be it an entrant or an existing firm – commits to innovating in that market. Then, the innovator learns whether she can improve upon the existing leading technology (with probability $1 - \beta$) or not (with probability β).

One possible interpretation is that the leader's patent may be infringed if the innovation builds upon the leading technology q_{it}, which happens with probability $\beta > 0$ (thus giving β an interpretation of patent rights enforcement). In case of a patent infringement, the incumbent firm would possess all bargaining power and could extract the entrant's future profits in its entirety. Anticipating this outcome, entrants optimally decide to build on the inferior public technology q_{-it} if and only if the leading technology is protected – and on the leading technology otherwise.

Hence, an entrant who has received a signal for market i builds upon the leading technology q_{it} with probability $(1 - \beta)$, and on the public technology q_{-it} with probability β.

The innovative steps λ achieved by entrants are drawn from a time-invariant distribution $F(\lambda)$. For tractability, I assume λ to be Pareto distributed, i.e. $F(\lambda) = 1 - \lambda^{-\alpha}$, where α denotes the Pareto shape parameter.

Innovation Success and Market Entry Whether an innovator can successfully enter a market depends on whether or not he can supersede the leading technology at the frontier. If an innovation builds directly on the leading technology, it is always successful since any innovation step satisfies $\lambda \geq 1$; I will refer to these as *frontier* innovations. On the other hand, if an innovation builds on the public technology, an innovator can successfully enter the market if $\lambda q_{-it} \geq q_{it}$.

In order to keep track of the market-specific technological gaps, let $\mu_{it} \equiv \frac{q_{it}}{q_{-it}}$ denote the technological gap in market i. Thus, if the leading technology is unavailable to the innovator, only innovative steps $\lambda \geq \mu$ are large enough to capture the market and worthwhile to be protected by a patent. Smaller innovative steps $\lambda < \frac{q_{it}}{q_{-it}}$ are not socially wasteful though: They push the public technology from q_{-it} to λq_{-it} and hence still contribute to the stock of technological knowledge. In either case, the lower of the two technological levels (i.e. min $\{\lambda q_{-it}, q_{it}\}$) becomes the new public technology, while max $\{\lambda q_{-it}, q_{it}\}$ constitutes the new frontier.

Labor Market Clearing Labor is used in production, for market search by incumbents, for R&D activities by incumbents, and by potential entrants (which, in turn, hire workers to do both market search and R&D). There is a common wage w_t that clears that labor market, and it holds at each instant t that

$$\underbrace{L_{P,t}}_{\text{production}} + \underbrace{L_{S,t}}_{\text{search}} + \underbrace{L_{R\&D,t}}_{\text{R\&D}} + \underbrace{L_{entry,t}}_{\text{entry}} = L, \tag{3}$$

where $L \equiv 1$.

2.1 Static Equilibrium

The static equilibrium takes the leading technology q_{it} and the public technology q_{-it} for each product line as given. Exploiting that the final good sector spends the same amount Y_t on each intermediate variety and that the leader of product line i sets the price equal to the marginal cost of the fringe firms, i.e. $p_{it} = w_t/q_{-it}$, the demand for variety i follows as

$$y_{it} = \frac{Y_t}{p_{it}} = \frac{q_{-it}}{w_t/Y_t} = \frac{q_{-it}}{\omega_t},$$

where $\omega_t \equiv w_t/Y_t$ denotes the wage normalized by total output. The labor share of the economy is equal to

$$\text{laborshare}_t = \frac{w_t L}{GDP_t}, \tag{4}$$

where $GDP_t = Y_t + w_t\big(L_{S,t}(\omega_t) + L_{R\&D,t}(\omega_t) + L_{entry,t}(\omega_t)\big)$ is defined as total final output plus R&D and search expenditures (by both incumbents and entrants). Since the labor supply L is normalized to unity and using $\omega_t \equiv w_t/Y_t$, it follows

$$\text{laborshare}_t = \frac{\omega_t}{1 + \omega_t\big(L_{S,t}(\omega_t) + L_{R\&D,t}(\omega_t) + L_{entry,t}(\omega_t)\big)}. \tag{5}$$

In equilibrium, the markups of active firms are equal to the productivity advantages over their respective competitors (fringe firms). Hence, the markup charged by the leader in line i at time t is

$$\mu_{it} = \frac{p_{it}}{mc_{it}} = \frac{q_{it}}{q_{-it}}.$$

Operating profits are thus given by

$$\pi(\mu_{it}) = (p_{it} - mc_{it})y_{it} = (1 - \frac{1}{\mu_{it}})Y_t \geq 0.$$

Labor demand in product line i follows as

$$l_{it} = \frac{y_{it}}{q_{it}} = \frac{1}{\mu_{it}\,\omega_t}. \tag{6}$$

Substituting this into the labor market clearing condition (3) yields

$$\int_0^1 \frac{1}{\mu_{it}\,\omega_t}\,di + L_{S,t}(\omega_t) + L_{R\&D,t}(\omega_t) + L_{entry,t}(\omega_t) = L.$$

Using $L_{P,t}(\omega_t) = L - \big(L_{S,t}(\omega_t) + L_{R\&D,t}(\omega_t) + L_{entry,t}(\omega_t)\big)$, the normalized wage is pinned down as

$$\omega_t = \frac{\int_0^1 \frac{1}{\mu_{it}}\,di}{L_{P,t}(\omega_t)}, \tag{7}$$

and the labor share follows from equation (5) as

$$\text{laborshare}_t = \frac{1}{\psi_t^h\, L_{P,t}(\omega_t) + \big(L_{S,t}(\omega_t) + L_{R\&D,t}(\omega_t) + L_{entry,t}(\omega_t)\big)} \leq 1, \tag{8}$$

where $\psi_t^h \equiv \frac{1}{\int_0^1 \frac{1}{\mu_{it}}\,di} \geq 1$ defines the harmonic mean of markups. Note that $\text{laborshare}_t < 1$ whenever some markups μ_{it} are strictly greater than 1. Given that the amount of labor used for production does not vary dramatically across different parametrizations, the key driver of the labor share is the average inverse markup. Since there is no capital in the model, the remaining part of GDP constitutes the profit share.

It is instructive to express aggregate output Y_t in terms of the production possibility frontier $Q_t L_{P,t}$ and a distortive term. Aggregate output is lower than $Q_t L_{P,t}$ if and only if there is heterogeneity in markups and thus in labor demands, since the marginal product of labor is then not equalised across varieties. Formally, using $y_{it} = q_{-it}/\omega_t = (q_{it}/\mu_{it})/\omega_t$ and equation (2), aggregate output can be written as

$$Y_t = \exp\Big(\int_0^1 \ln \frac{q_{it}}{\mu_{it}\,\omega_t}\,di\Big) = \frac{Q_t}{\psi_t^g\,\omega_t} = \frac{Q_t L_{P,t}}{\psi_t^g \int_0^1 \frac{1}{\mu_{it}}\,di}, \tag{9}$$

where $1/(\psi_t^g \int_0^1 \frac{1}{\mu_{it}}\, di) \leq 1$ is the static macroeconomic distortion to total output due to heterogeneous markups[3] and $\psi_t^g \equiv \exp(\int_0^1 \ln \mu_{it}\, di)$ is the geometric average of product-level markups. In the absence of heterogeneous markups, i.e. if $\mu_{it} = \bar{\mu}_t$, the geometric and the arithmetic averages coincide and the economy produces at the production possibility frontier without distortion.

2.2 Balanced Growth Path

I now characterize the balanced growth path of the economy. Time indexes are usually omitted for convenience, and I sometimes use asterisks to emphasize steady state variables. In steady state, the Euler equation (1) reads

$$r^* = \rho + g^*, \tag{10}$$

where g^* is the steady state growth rate of consumption C_t. Aggregate output Y_t and the wage w_t also grow at the common growth rate g^*.

Next, let us consider the distribution of markups, which is the key object to determine the labor share and factor misallocation but – by affecting profits – also shapes the incentives to innovate. On the balanced growth path, the distribution of markups μ is stationary, and factor shares are constant.

Productivity gap distribution The distribution of interest is the stationary distribution of productivity gaps denoted by $H(\mu)$, which is is an equilibrium outcome as it depends on s^*. The following key theoretical result can be established:

Proposition 1. *Productivity gaps are Pareto distributed with an endogenously determined shape parameter $\frac{\alpha}{s^*}$, that is, $H(\mu) = 1 - \mu^{-\frac{\alpha}{s^*}}$,*

where s^* is the optimal search intensity of both entrants and incumbents on the balanced growth path. [detailed proof of the proposition will follow here]

As a corollary, it follows that the distribution of targeted markets is given by

$$G_s(\mu) = 1 - (1 - H(\mu))^s. \tag{11}$$

Proof: Innovators always target the market with the lowest markup in the sample. Given a search intensity (and hence a sample size) of s^*, the cumulative distribution function of the minimum of s^* iid draws is $1 - (1 - H(\mu))^{s^*}$.

3. Using a slightly different notation, Peters (2019) refers to the markup distortion term $1/(\omega_t\, \psi_t) \leq 1$ as the *markup wedge*.

The intuition behind this is as follows: The targeted productivity gap is the minimum of s i.i.d. draws from the cross-sectional distribution $H(\mu)$.

Combining Proposition (1) and equation (11), it follows that the productivity gaps that can be targeted when screening N markets is

$$G_s(\mu) = 1 - \mu^{-\alpha \frac{s}{s^*}}, \tag{12}$$

where s^* is taken as given but s is a choice variable ($s = s^*$ holds on the unique balanced growth path). By searching more markets, a prospective innovator tends to find a market with a currently less dominant (i.e. lower markup) leader. Thus, the distribution of markets that can be targeted by an innovator with search intensity s_2 is first-order stochastically dominated by the distribution of markets that can be targeted by an innovator with a lower search intensity $s_1 < s_2$.[4]

Entrant Behavior The entrant problem determines the optimal number of markets s^* that are screened simultaneously in the exploration stage, as well the total mass of prospective entrants $\mathcal{M}_0$ that are actively exploring markets at any point in time on the balanced growth path. The search intensity s^* is chosen to maximize expected profits, whereas $\mathcal{M}_0$ ensures that the free entry condition is satisfied in equilibrium. The equilibrium is unique irrespective of whether there is an entry congestion force ($\delta > 0$) or not: While the cost of entry is weakly increasing in the number of entrants $\mathcal{M}_0$, the value of a new firm is strictly decreasing in $\mathcal{M}_0$ because of intensified creative destruction.

The likelihood of an innovation being successful – that is, the likelihood of overtaking the current leader – depends on the steady state distribution of productivity gaps and on s. By searching more intensely, innovators can increase their *entry success* rate. However, an innovator cannot affect the productivity advantage in case the innovation effort succeeds: conditional on superseding the leader's technology, the new productivity gap will be Pareto distributed, irrespective of the current cross-sectional distribution of productivity gaps[5].

The *entrant success* probability conditional on innovating upon the *public* technology equals the probability that the randomly drawn innovation exceeds the current productivity gap and thus is equal to $(1 - F(\mu))$. The optimization problem of a prospective

4. Formally, for any $s_2 > s_1 \geq 1$, it holds that $G_{N_2}(\mu) < G_{N_1}(\mu) \forall \mu$.

5. Formally: Suppose $\lambda \sim \text{Pareto}(\alpha)$. Then for any constant $\zeta \geq 1$, conditional on $\lambda \geq \zeta$, it holds $(\frac{\lambda}{\zeta}) \sim \text{Pareto}(\alpha)$, i.e. the Pareto distribution is preserved.

entrant is then given by

$$\max_{s} \underbrace{\left[1-\beta+\beta\times\int_{1}^{\infty}(1-F(\mu))\,\mathrm{d}G_{s}(\mu)\right]}_{\text{success probability, }\uparrow\text{ in }s}\times E\left(V_{\text{entrant,t}}\right) - \frac{1}{\theta}(\mathcal{M}_0)^{\delta} - \underbrace{\frac{1}{\phi}\eta s^{\frac{1}{\eta}} w_t}_{\text{exploration cost}}, \quad (13)$$

where $E\left(V_{\text{entrant,t}}\right)$ is the expected value of a firm entering at time t. Since each entrant is atomistic, $E\left(V_{\text{entrant,t}}\right)$ does not depend on s directly, hence the key tradeoff is between the entrant success probability (first term) and the exploration costs. Moreover, the free entry condition sets equation (13) equal to zero, implicitly pinning down the mass of potential entrants $\mathcal{M}_0$.

Firm Problem Each incumbent firm chooses its search intensity s and innovation rate x. The firm problem is given by

$$\max_{x,s}\ \underbrace{x\left[1-\beta+\beta\times\int_{1}^{\infty}(1-F(\mu))\,dG_{s}(\mu)\right]}_{\text{success probability, }\uparrow\text{ in }s}\times E\left(V_{t}\right) - \frac{1}{\bar{\theta}}\gamma x^{1/\gamma} w_t - \underbrace{x\frac{1}{\phi}\eta s^{\frac{1}{\eta}} w_t}_{\text{exploration cost}}, \quad (14)$$

where $E\left(V_{t}\right)$ denotes the expected value of an additional product for an incumbent. It can be shown that the value of a firm is additive in its individual product values, thus it holds that $E\left(V_{t}\right)=E\left(V_{\text{entrant,t}}\right)$. Let $E(v_t)\equiv\frac{E(V_t)}{w_t}$ denote the expected product value normalized by the wage rate w_t. Note that it again holds $E\left(v_{\text{entrant,t}}\right)=E\left(v_{t}\right)$. Making use of the fact that on the balanced growth path, the value function $V_t(\mu)$ grows at the same rate as w_t, it follows that $E\left(v_t\right)$ is constant. Exploiting equation (11) and setting the first order condition of the firm problem equation (14) with respect to s to zero thus yields

$$x\frac{s^*}{(s+s^*)^2}\beta E(v) = x\frac{1}{\phi}s^{1/\eta-1}.$$

Similarly, setting the first order condition of the entrant problem equation (13) with respect to s to zero implies again

$$\frac{s^*}{(s+s^*)^2}\beta E(v) = \frac{1}{\phi}s^{1/\eta-1}.$$

Thus, all firms and entrants search with the same intensity s^* and exploiting $s=s^*$ yields

$$s^* = \left(\frac{\beta\phi E(v)}{4}\right)^{\eta}. \quad (15)$$

Substituting s^* from equation (15) into (13) and setting this equation equal to zero yields the free entry condition

$$\left(1 - \frac{\beta}{2}\right) \times E(v_{\text{entrant}}(\mathcal{M}_0)) \times \left(1 - \frac{\eta}{\frac{4}{\beta} - 2}\right) - \frac{1}{\theta}(\mathcal{M}_0)^{\delta} = 0, \tag{16}$$

where the notation highlights that the firm values depend on the total mass $\mathcal{M}_0$ of prospective entrants that are actively exploring markets. Higher $\mathcal{M}_0$ implies a higher rate of creative destruction, thus decreasing the expected lifetime and value of firms. At the same time, more entry increases the effective R&D cost of entrants on the right-hand-side of equation (16). On the balanced growth path, $\mathcal{M}_0$ is constant and uniquely pinned down. The expected value of an entrant firm (normalized by the wage w_t) follows from the free entry condition as

$$E\left(v_{\text{entrant}}\right) = \frac{4}{4 - 2\beta - \eta\beta} \times \frac{1}{\theta}(\mathcal{M}_0)^{\delta}. \tag{17}$$

Substituting this back into equation (15) leads to the following closed-form solution for s^*:

$$s^* = \left(\frac{\frac{\phi}{\theta}(\mathcal{M}_0)^{\delta}}{\frac{4}{\beta} - 2 - \eta}\right)^{\eta}. \tag{18}$$

The importance of the optimal search intensity s^* is manifold: it determines the distribution of markups, the average markup, factor shares and factor misallocation across firms. Fundamentally, s^* determines the shape of the productivity gap distribution (and thus the markup distribution) as is directly implied by equation (1), so it is worthwhile to discuss how the different parameters affect s^*. Given that we live in an environment where ideas are getting harder to find (Bloom et al. 2020), changes in $\theta \downarrow, \tilde{\theta} \downarrow$ are naturally the most relevant comparative statics. As the research productivity declines, incentives to search markets increase:

$$\frac{\partial s^*}{\partial \theta} < 0. \tag{19}$$

Note that while the entrant research productivity θ enters directly, incumbent research productivity $\bar{\theta}$ has an indirect effect through $\mathcal{M}_0$. A lower incumbent research productivity $\bar{\theta}$ leads ceteris paribus to less creative destruction and thus makes entry more attractive, thereby increasing $\mathcal{M}_0$ and s^*. Nevertheless, $\mathcal{M}_0$ declines when ideas are getting harder to find, because lower entrant research productivity disincentivizes entry sufficiently.

Moreover, a higher probability to build on the public technology β implies a higher s^* because it becomes more important to avoid markets with dominant (high productivity gap) leaders:

$$\frac{\partial s^*}{\partial \beta} > 0 \tag{20}$$

Furthermore, higher η means that exploration costs are less responsive to s, hence s^* is more sensitive to parameter changes if η is high. Further, s^* depends positively on the ratio $\frac{\phi}{\theta}$ of the cost parameters of entrant's R&D relative to market exploration.

Returning to equation (14) and setting the first order condition with respect to x to zero yields

$$(1-\beta/2)E(V_t) = \frac{1}{\tilde{\theta}}x^{(1-\gamma)/\gamma}w_t + \frac{1}{\phi}\eta s^{1/\eta}w_t, \tag{21}$$

and again using the same notation as before, this can be written compactly as

$$(1-\beta/2)E(v) = \frac{1}{\tilde{\theta}}x^{(1-\gamma)/\gamma} + \frac{1}{\phi}\eta s^{1/\eta}. \tag{22}$$

Substituting for s^* and for $E(v)$ and rewriting yields

$$(4/\beta-2)\frac{1}{4/\beta-2-\eta} \times \frac{1}{\theta}(\mathcal{M}_0)^{\delta} = \frac{1}{\tilde{\theta}}x^{(1-\gamma)/\gamma} + \eta\left(\frac{\frac{1}{\theta}(\mathcal{M}_0)^{\delta}}{\frac{4}{\beta}-2-\eta}\right), \tag{23}$$

which can be further simplified to

$$\frac{1}{\theta}(\mathcal{M}_0)^{\delta} = \frac{1}{\tilde{\theta}}x^{(1-\gamma)/\gamma}, \tag{24}$$

Growth rate It can be shown that the steady state growth rate of the economy is equal to the product of three terms,

$$g^* = \big(x^* + \mathcal{M}_0\big) \times \big(1-\frac{\beta}{2}\big) \times E\big(\log(\mu_{new})\big). \tag{25}$$

The intuition is as follows. The term $(x^* + \mathcal{M}_0)$ is the rate at which the economy generates innovations; $(1-\frac{\beta}{2})$ is the probability that innovations are successful (in the sense of capturing the market); and $E(\log(\mu_{new}))$ is the average logarithm of productivity gaps of a new cohort of firms, indicating how much the technological frontier is pushed forward. Using the Pareto property that new cohorts of productivity gaps are identically distributed as the underlying draws of innovative steps (and thus $\sim$ Pareto(α)), it can be shown that $E(\log(\mu_{new})) = 1/\alpha$. Under the additional assumption that each successful innovation corresponds to the entry of a new establishment in the data, the first two terms are equal to the establishment entry rate. Put differently, the following relationship can be established:

$$\frac{g^*}{\text{EstabEntryRate}} = E\big(\log(\mu_{new})\big) = \frac{1}{\alpha} \tag{26}$$

Firm Values Normalized firm values $v(\mu) \equiv \frac{V_t(\mu)}{w_t}$ satisfy

$$\rho v(\mu) = \underbrace{\Big(1-\frac{1}{\mu}\Big)\frac{Y}{w} + \Big(\frac{1}{\gamma}-1\Big)\frac{1}{\tilde{\bar{\theta}}}\gamma(\tilde{x})^{\frac{1}{\gamma}}}_{\substack{\text{instantaneous payoff:}\\ \text{profits + option value of innovation}}} \tag{27}$$

$$- \underbrace{(\tilde{x}+\mathcal{M}_0)}_{\substack{\text{aggregate \#}\\ \text{of innovations}}} \underbrace{\frac{g_{s^*}(\mu)}{h(\mu)}}_{\substack{\text{gap-specific}\\ \text{correction of}\\ \text{hazard rate}\\ \text{(search effect)}}} \Bigg(v(\mu) - \underbrace{\beta \Pr(\lambda \le \mu)\, \mathbb{E}\Big(v\big(\frac{\mu}{\lambda}\big)|\lambda \le \mu\Big)}_{\text{leader survives but gap shrinks}}\Bigg), \tag{28}$$

The first term on the right-hand-side are the instantaneous profits of a firm with productivity gap μ, and the second term is the option value of doing R&D that the additional product adds to the firm. The hazard rate of an innovation being targeted towards a market with gap μ is given by $\frac{g_{s^*}(\mu)}{h(\mu)}(\tilde{x}+\mathcal{M}_0)$, where $(\tilde{x}+\mathcal{M}_0)$ is the innovation arrival rate of the economy and $\frac{g_{s^*}(\mu)}{h(\mu)}$ reflects the fact that high productivity gap markets are less likely to be targeted by innovators. It holds that $g_{s^*}(\mu) = \alpha\mu^{-\alpha-1}$ and $h(\mu) = \frac{\alpha}{s^*}\mu^{-\frac{\alpha}{s^*}-1}$.

Given that $v(1) = 0$, firms values according to equation (27) can be computed sequentially since $v(\mu)$ only depends on firm values $v(\tilde{\mu})$ where $\tilde{\mu} < \mu$.

Thus, the steady state of the model can be solved for as follows: $\mathcal{M}_0$ and x are such that equation (24) holds, and the 1-product firm value function defined in equation (27) implies an expected entrant firm value $E(v_{\text{entrant}}(\mathcal{M}_0))$ that is consistent with the free entry condition (16).

Results on balanced growth path

In steady state, it can be shown that for entrants, the ratio of exploration costs relative to R&D expenditures at the firm level is given by

$$\frac{1}{\gamma}\frac{\eta}{4}\left(\frac{1}{1/\beta - 1/2 - \eta/4}\right), \tag{29}$$

and for entrants the ratio is

$$\frac{\eta}{4}\left(\frac{1}{1/\beta - 1/2 - \eta/4}\right). \tag{30}$$

In both cases, the ratio of exploration costs relative to R&D expenditures is constant across steady states (when research productivities θ and $\bar{\theta}$ decline). This again illustrates that as ideas get harder to find and R&D becomes costlier, there is more market search for any given innovation.

Using proposition (1), the cost-weighted average markup can be simplified to

$$\text{average markup} = \frac{\alpha + s^*}{\alpha}, \tag{31}$$

which is increasing in s^*. The intuition is as follows: More market search implies more selection based on market power, and hence higher markups on average.

Similarly, the labor share of the economy follows from equation (8) as

$$\text{laborshare} = \frac{1}{\frac{\alpha}{\alpha+s^*} L_{P,t}(\omega_t) + \big(L_{S,t}(\omega_t) + L_{R\&D,t}(\omega_t) + L_{entry,t}(\omega_t)\big)} \leq 1, \tag{32}$$

and the profit share follows as $1 -$ laborshare.

3 Predictions when ideas get harder to find

Research productivity has declined at a rate of up to five percent per annum over the last decades (Bloom et al. 2020). To assess the ramifications of these trends, I decrease the research productivity parameters of incumbents and entrants by the same factor while leaving the search productivity and all other parameters of the model unchanged.

While this section provides the intuition of the economic forces at work, section 5 quantifies the model predictions over the course of time. I allow for a recalibration of the size of the effective labor force in order to capture population growth and higher labor force participation as well as a market size effect due to international trade; given that the increase in the effective labor force is outweighed by the decline in research productivity, the former can mitigate but not prevent falling growth.

As innovations become effectively more costly, the incentives to search markets increase, thereby leading to stronger selection on profitability, higher and more dispersed markups, and a higher profit share. Comparing the new to the initial steady state, there is a direct effect of lower research productivity depressing the demand for R&D workers; however, this effect is dominated by an endogenous effect of a higher profit share that incentivizes R&D, thus raising the population share of researchers across steady states.

The model predicts a productivity growth slowdown since the increase in researchers is not enough to compensate for the decline in research productivity. Finally, the aggregate innovation rates by entrants and incumbents both decrease (leading to generally lower business dynamism and older firms), but the entrants' innovation rate decreases relatively more because the entry margin is more elastic (which holds under very weak conditions). As a result, the firm entry rate declines and the average firm size increases. All these model predictions are in line with the empirical trends of the US economy over the last decades.

It is worthwhile to point out the implications of strengthened defensive capabilities in the spirit of Akcigit and Ates (2021). If the degree of barriers to frontier innovation were increased, market search would be further incentivized, because overcoming large technological gaps becomes disproportionally harder. Markups become more dispersed and market power increases, mirroring the data trends. Thus, strengthened defensive behavior by industry leaders can qualitatively account for the macroeconomic experience of the US economy over the last decades. However, the model predicts changes that are quantitatively too small in comparison to the data. Moreover, the existing empirical evidence for a trend in defensive capabilities of incumbents is sparse, because it is an abstract parameter that is inherently difficult to measure and tightly linked to institutional details. On the other hand, declining research productivity is an ongoing technological process. Nevertheless, it is plausible that an increased role of defensive behavior contributed to the macroeconomic trends, which is a view corroborated by Akcigit and Ates (2021).

4 Calibration

In this section, I calibrate the model to the pre-1980 economy; this will serve as a point of comparison in the next section when I study the shift towards a new steady state of the 2010s.

There are ten parameters: $\{L, \beta, \rho, \alpha, \eta, \gamma, \phi, \theta, \tilde{\theta}, \delta\}$. The inelastic labor supply is normalized to unity, $L = 1$. I set the discount rate $\rho = 1\%$ implying a discount factor of 99 percent. Following Acemoglu et al. (2018), the R&D elasticity is equal to $\gamma = 0.5$, which implies a quadratic R&D costs specification. Since the search technology of my model is a methodological novelty, there exist no comparable estimates in the literature; I set the R&D elasticity equal to the search elasticity so that $\eta = \gamma = 0.5$. I choose the intermediate value $\beta = 0.5$ for the degree of barriers to frontier innovation, meaning that innovators build upon the leader or upon the follower with equal probability. Lastly, I assign a congestion externality for entrants of $\delta = 0.7$. The remaining four parameters are internally calibrated by exactly matching the following four moments:

Average TFP growth of the US economy (1.82% during the 30-year-period 1949-1978 according to the Fernald series) is informative of the incumbent research productivity $\tilde{\theta}$. Entrant research productivity θ can then be set to match the contribution of entrants to TFP growth; I target a value of 25% that is equal to the contribution of entrants to TFP growth among U.S. manufacturing establishments between 1977 and 1982 (Bartelsman and Doms 2000). The key moment to calibrate the Pareto shape of innovations, α, is the firm entry rate; a value of 0.13 is taken from the Census Bureau's Business Dynamics Statistics for the year 1980. Finally, the search productivity parameter can be set by targetting the average markup, which I set to 1.1 corresponding to the simple average between the estimated values of i) De Loecker, Eeckhout, and Unger (2020), ii) Eggertsson, Robbins, and Wold (2021), and iii) Nekarda and Ramey (2020).

The lower panel of table 2 shows the internally calibrated parameters. While the productivity parameters cannot be directly interpreted, a brief discussion of the Pareto shape parameter α is in order: The implied value of 13.2 is so high because TFP growth is less than two percent while the firm entry rate (with each firm entry corresponding to an instance of creative destruction) is thirteen percent – thus, the average innovative step has to be small and the Pareto tail index relatively high.

5 Macroeconomic Trends since the 1980s

Since the 1980s, the US economy has experienced a number of salient macroeconomic trends. These include i) the rise of market power (De Loecker, Eeckhout, and Unger 2020); ii) a decrease in the firm entry rate (R. A. Decker et al. 2016); iii) a declining share of activity from firms of at most five years; iv) an increase in average firm size (Hopenhayn, Neira, and Singhania 2018); and v) a growth slowdown (interrupted by IT-fueled boom in late 1990s) despite a higher R&D-to-GDP ratio. Trends ii), iii) and iv) are symptoms of what is commonly referred to as declining business dynamism (Akcigit and Ates 2021).

Through the lens of my model, the decline in research productivity (*ideas are getting harder to find* à la Bloom et al. 2020) can jointly explain the other macroeconomic trends. As established in proposition 1, the key object for the distribution of markups is the search intensity s^*. I take the following stand on the parameters: The search elasticity η and the entry congestion δ are fundamental time-invariant parameters. Of the remaining three parameters, either i) a decrease in the research productivities θ (for entrants) and $\tilde{\theta}$

Assigned	Value	Description
ρ	0.01	Discount rate
β	0.5	Innovation barriers
η	0.5	Search elasticity
γ	0.5	Incumbent R&D elasticity
δ	0.7	Entry externality
L	1	Size of labor force

Estimated	Value	Description	Key moment	Model	Data
$\tilde{\theta}$	0.63	**Incumbent R&D prod.**	TFP growth	0.0182	0.0182
θ	0.45	**Entrant R&D prod.**	Entrants' % growth	0.25	0.25
α	13.2	Pareto shape innov.	Firm entry rate	0.13	0.13
ϕ	25.2	Search productivity	Average markup	1.1	1.1

Table 1: Baseline parameter values

(for incumbents), ii) an increase in the search productivity ϕ, or iii) an increase in barriers to frontier innovation β induce a rising search intensity s^* and, consequently, higher markups.

Both the search productivity and the barriers to frontier innovation are abstract parameters and difficult (if not impossible) to measure. An increase in barriers to frontier innovation would be reminiscent of the explanation provided by Akcigit & Ates (2021), who use a theoretical framework to find that declining knowledge diffusion from leaders to laggards is consistent with the macroeconomic trends. The decline in knowledge diffusion is most likely a result of intensified defensive behavior by leaders, but it is difficult to measure this empirically. My results corroborate their findings, but importantly, I show that there is an alternative explanation that is better supported by existing evidence: *ideas are getting harder to find*.

	pre-1980 s.s.	2010 s.s.	Change: Model	Change: Data	**Model/Data**
Targeted moments					
TFP growth	0.0182	0.0096	-47%	-72%	**65**%
Entrants' % growth	0.25	0.206	-18%	-	-
Firm entry rate	0.13	0.064	-51%	-39%	**131**%
Average markup	1.1	1.15	$+4.5\%$	$+7\%$	**64**%
Untargeted moments					
Average firm size	2.16	2.44	$+13\%$	$+15\%$	**87**%
Profit share	0.041	0.084	$+105\%$	$+75\%$	**140**%
R&D workers % of L	0.076	0.086	$+13\%$	$+50\%$	**26**%

Table 2: Steady state comparison after a 75% decline in research productivities

Bloom et al. (2020) document a general decline in research productivity that is not specific to any technology or industry. While the study by Bloom et al. (2020) is entirely empirical, I show that their finding can, through the lens of my model, account for all other macroeconomic trends jointly. To illustrate these macroeconomic ramifications, I decrease the research productivities θ (for entrants) and $\tilde{\theta}$ (for incumbents) by 75% relative to the initial steady state (corresponding to a 5% per annum decline over 30 years).

Table 2 summarizes the results and compares the model predictions to the data counterparts, where the data moments for the 2010-economy are computed in the same way and based on the same data sources as in the initial calibration. For interpretability, I focus on relative changes in all moments. The model predicts a five percentage point increase in the average markup (corresponding to a 4.5% increase in relative terms), thereby explaining as much as two thirds of the overall rise of market power. Thus, my theory connects a substantial share of the rise of market power (De Loecker, Eeckhout, and Unger 2020) to the decline in research productivity (Bloom et al. 2020).

Moreover, while TFP growth decreased by 72% in the data, the model can account for a 47% decrease (from 1.82% to 0.96% per annum) and thus for roughly two thirds of the decline. The model slightly overestimates the response of the firm entry rate (-51% relative to -39% in the data) and underestimates the response of the average markup ($+4.5\%$ relative to $+7\%$ in the data), but all predictions have the correct sign and are of the right order of magnitude. The lower panel of table 2 shows that the model's predictions are also in line with the data for a number of untargeted moments.

6 Conclusions

This is the first paper that connects the rise of market power to the decline in research productivity. To this effect, I build a novel theory of directed innovation where innovators conduct market search before targeting a specific market and undertaking R&D efforts. In an environment of declining research productivity (as documented by Bloom et al. 2020), R&D becomes costlier; as a consequence, innovators are more inclined to spend on market search in order to avoid wasteful R&D expenditures in markets where they would struggle to compete with superior incumbents. These superior (high-markup) incumbents, in turn, are now less likely to be targeted by innovators and thus more likely to survive than they were four decades ago. The model predicts an increase of the economy-wide average markup in the US from 1.1 in the 1980s to 1.15 in the 2010s, which is in the ballpark of estimates known from the literature. Similarly, the model mirrors other relevant macro trends both qualitatively and quantitatively.

Evidently, declining research productivity has a lot of explanatory power. If research productivity will decline in the future, the model would predict a continuation of the ongoing trends – and thus a further rise of markups and firm sizes, and a further fall of the labor share, the firm entry rate and aggregate productivity growth.

References

Acemoglu, D., Akcigit, U., Alp, H., Bloom, N., and Kerr, W. 2018. "Innovation, Reallocation, and Growth." *American Economic Review* 108 (11): 3450–3491.

Aghion, P., Bergeaud, A., Boppart, T., Klenow, P. J., and Li, H. 2021. *A Theory of Falling Growth and Rising Rents.* Working Paper.

Aghion, P., Harris, C., and Vickers, J. 1997. "Competition and growth with step-by-step innovation: An example." *European Economic Review* 41 (3–5): 771–782.

Aghion, P. and Howitt, P. 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60 (2): 323–351.

Akcigit, U. and Ates, S. T. 2021. "Ten Facts on Declining Business Dynamism and Lessons from Endogenous Growth Theory." *American Economic Journal: Macroeconomics* 13, no. 1 (January): 257–98.

Argente, D., Baslandze, S., Hanley, D., and Moreira, S. 2020. *Patents to Products: Product Innovation and Firm Dynamics.* Working Paper. Fed Atlanta Working Paper Series.

Arrow, K. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors,* 609–626. NBER Chapters. National Bureau of Economic Research, Inc, October.

Bartelsman, E. J. and Doms, M. 2000. "Understanding Productivity: Lessons from Longitudinal Microdata." *Journal of Economic Literature* 38 (3): 569–594.

Bloom, N., Jones, C. I., Van Reenen, J., and Webb, M. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review* 110, no. 4 (April): 1104–44.

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

De Ridder, M. 2020. *Market Power and Innovation in the Intangible Economy.* Working Paper.

Decker, R., Haltiwanger, J., Jarmin, R., and Miranda, J. 2014. "The Role of Entrepreneurship in US Job Creation and Economic Dynamism." *Journal of Economic Perspectives* 28, no. 3 (September): 3–24.

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

Eggertsson, G. B., Robbins, J. A., and Wold, E. G. 2021. "Kaldor and Piketty's facts: The rise of monopoly power in the United States." *Journal of Monetary Economics* 124 (S): 19–38.

Elsby, M., Hobijn, B., and Sahin, A. 2013. "The Decline of the U.S. Labor Share." *Brookings Papers on Economic Activity* 44 (2): 1–63.

Galasso, A. and Schankerman, M. 2015. "Patents and Cumulative Innovation: Causal Evidence from the Courts." *The Quarterly Journal of Economics* 130 (1): 317–369.

Grossman, G. M. and Helpman, E. 1991. "Quality Ladders in the Theory of Growth." *Review of Economic Studies* 58 (1): 43–61.

Hopenhayn, H., Neira, J., and Singhania, R. 2018. *From Population Growth to Firm Demographics: Implications for Concentration, Entrepreneurship and the Labor Share.* Working Paper, Working Paper Series 25382. National Bureau of Economic Research.

Klette, T. J. and Kortum, S. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy* 112 (5): 986–1018.

Liu, E., Mian, A., and Sufi, A. 2019. *Low Interest Rates, Market Power, and Productivity Growth.* Working Paper 25505. National Bureau of Economic Research.

Manera, A. 2021. *Competing for Inventors: Market Concentration and the Misallocation of Innovative Talent.* Working Paper, Working Paper Series.

Nekarda, C. J. and Ramey, V. A. 2020. "The Cyclical Behavior of the Price-Cost Markup." *Journal of Money, Credit and Banking* 52 (S2): 319–353.

Nelson, R. R. 1959. "The Simple Economics of Basic Scientific Research." *Journal of Political Economy* 67 (3): 297–306.

Olmstead-Rumsey, J. 2022. *Market Concentration and the Productivity Slowdown,* Working Paper. London School of Economics.

Peters, M. 2019. *Heterogeneous Markups, Growth and Endogenous Misallocation.* Working Paper. Yale University.

Peters, M. and Walsh, C. 2021. *Population Growth and Firm Dynamics.* Working Paper.