Endogenous Growth and Rising Obsolescence

By Yoshiki Ando (BU, TPRI) and James Bessen (BU, TPRI)

Preliminary Version
Please do not share without permission.

Abstract:
Recent evidence suggests that obsolescence rates have risen for R&D (Ando, Bessen, and Wang 2025). How does rising obsolescence affect productivity growth? This paper presents a growth model where technological opportunities of varied quality arrive exogenously for both new products and replacement products. Researchers choose which opportunities are feasible to develop. But the expected rents are greater for replacement innovations than for new product innovations of comparable quality. The average replacement innovation has lower quality than the average new product innovation, so it contributes less to aggregate productivity. This means that a *shift* toward more replacement innovations (higher obsolescence) can reduce productivity growth substantially for a significant period.

This work was supported in part through a grant from the Alfred P. Sloan Foundation (Grant #2023-21111).

Introduction

Since 1980, real R&D spending in the US has risen 50% faster than GDP, more than quadrupling.[1] Nevertheless, the rate of US productivity growth has not increased and appears to have declined somewhat.[2] Several papers have argued that research productivity is declining and some argue that productivity growth is not increasing because "ideas are getting harder to find" (Bloom et al. 2020; Boeing and Hünermund 2020; Griliches 2009; S. Kortum 1993; Caballero and Jaffe 1993; Segerstrom 1998).

But a comprehensive analysis of US firms finds that since 1976, the gross marginal returns to R&D have risen sharply as well as the output elasticity of R&D (Ando, Bessen, and Wang 2025). Moreover, rising returns to R&D are consistent with rising R&D spending, while declining research productivity seems at odds with greater spending. But rising gross marginal returns accompanied by stagnant aggregate productivity growth does pose a puzzle.

This paper presents a mechanism that can reconcile growing R&D with stagnant aggregate productivity growth, at least for a period of time. This mechanism arises when R&D investment heightens technological rivalry between firms, leading to more rapid replacement of one innovation by another. A corollary of rising R&D effectiveness is rising effectiveness against rivals, leading to more rapid replacement. This change constitutes a shift in the *direction* of innovation that may have effects on productivity growth independently of the rate of innovation.

More rapid replacement means that the obsolescence rate rises. Indeed, there is evidence of rising R&D obsolescence. One piece of evidence comes from the rise in marginal returns to R&D noted above. In equilibrium, the gross marginal return to R&D equals the user cost of the R&D capital stock. Ignoring capital gains in R&D (because of limited resale), the user cost is the real after-tax price of R&D times the sum of the real interest rate plus the R&D depreciation/obsolescence rate. Since 1980, the after-tax price and interest rate have not increased or have decreased. By a process of elimination, the large increase in the returns to R&D implies a substantial increase in obsolescence rates in

[1] BEA Series Y694RX. These figures do not include software developed for internal use which has risen even more sharply.

[2] University of Groningen and University of California, Davis, Total Factor Productivity at Constant National Prices for United States [RTFPNAUSA632NRUG], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/RTFPNAUSA632NRUG, February 11, 2025. From 1983 to 2001, total factor productivity for the US grew 0.8% per year; from 2002 to 2019 growth was 0.6% per annum.

equilibrium.[3] Also, Ando, Bessen, and Wang (2025) estimate R&D obsolescence rates by measuring the impact of rivals' R&D; they estimate lower-bound obsolescence rates that have more than doubled.

Rising obsolescence means that relatively more R&D is directed toward replacing existing technologies, that is, the direction of R&D changes. The literature has long recognized differences in innovations between breakthrough and routine innovations, or between macro- and micro-innovations (Mokyr 1992), or between exploration and exploitation (March 1991). Our model makes a distinction between new products and replacement innovations: new products do not substitute for existing products while replacement improvements do. While this treatment is highly stylized—in reality new products have different degrees of substitutability with old ones—it highlights a key economic difference: the incentives to develop replacement innovations differ from the incentives to develop new product innovations. At least since Arrow (1962), replacement has been understood to alter innovation incentives. Aghion and Howitt (1992) incorporate this difference in their growth model via "business stealing": some or all of the quasi-rents formerly earned by an incumbent firm subsequently accrue to the firm with the replacement innovation.

We develop a model that borrows from existing growth models but includes both new products (as in Romer 1990) and replacement products (as in Grossman and Helpman 1991; Aghion and Howitt 1992). In each period, researchers are presented with technological opportunities for innovation, some for new products, and some for replacement products. These technological opportunities differ in quality where greater quality generates greater productivity and also greater quasi-rents for the researcher. The researchers endogenously choose to develop some of these opportunities—specifically, opportunities where the expected rents exceed a feasibility threshold determined by the researchers' alternative opportunities in production. We assume that both types of opportunities are drawn from a common quality distribution. But because of business stealing, a replacement opportunity offers greater incentive to the researcher than does a new product opportunity of the same quality. This means that researchers will face a lower feasibility threshold for replacement

[3] And evidence suggests that frictions such as monopsony or adjustment costs cannot account for rising disequilibrium.

innovations and, hence, the set of replacement innovations that they choose to develop will have lower average quality than the set of new product innovations. Lower average quality implies that replacement innovations contribute less to productivity growth on average and a rise in the replacement share of innovations (a rise in obsolescence) will reduce aggregate productivity growth. Although replacement innovations unambiguously increase the level of productivity, a rising replacement share reduces the rate of productivity growth, at least temporarily. We initially obtain this result assuming an exogenous increase in the replacement share of innovations. We then explore long-term changes that might affect the relative technological opportunities for replacement and new product development, including "fishing out" and the emergence of new technology fields.

Our paper takes a similar approach to Olsson (2005) in modeling technological opportunity in the long run, but we additionally consider the choice between new products and improvements. Several other papers also consider both improvements and new products (Howitt 1999; Jones and Williams 2000; Li 2000); our paper specifically considers the role of endogenous choice of technological opportunities. These papers, like ours, model the choice of individual researchers, abstracting away from the boundaries of the firm. Several other papers consider both replacement and new product innovations in the context of incumbent and entrant firms (Acemoglu et al. 2018; Akcigit and Kerr 2018; Garcia-Macia, Hsieh, and Klenow 2019). We abstract away from firm type to focus on differences that affect the rate of replacement.[4]

A variety of other explanations have been offered for the productivity slowdown including interest rates (Liu, Mian, and Sufi 2022), IT (Aghion et al. 2023), intangibles generally (Ridder 2024), declining research productivity (Bloom et al. 2020).

Model

Production

There are three sectors consisting of 1) a competitive final goods sector, 2) a measure N of monopoly producers of intermediate goods, each producing a single good, and 3) L_n

[4] Because we ignore the uncertainty of innovation, what we describe as a license between a researcher and an intermediate producer could equivalently be treated as a producer investing in its own R&D or contracting out R&D.

researchers developing new designs; some designs are for new products, some are for improvements on existing products. Each intermediate producer purchases an exclusive license to an infinitely lived patent for the design of their good from a researcher. Output, Y, is produced from production labor, L_y, and intermediate goods of quantities x_i and quality q_i in an additively separable production function,

$$Y = L_y^{\alpha} \int_0^N q_i^{\alpha} x_i^{1-\alpha} \, di, \quad 0 < \alpha < 1. \tag{1}$$

Total labor, $L = L_y + L_n$, is fixed.

The final goods producer chooses quantities x_i based on the prices p_i charged by the intermediate goods producers to maximize its profit

$$\pi^Y = \max_{x_i} \left[L_y^{\alpha} \int_0^N q_i^{\alpha} x_i^{1-\alpha} \, di - \int_0^N p_i x_i \, di - w L_y \right]. \tag{2}$$

The price of the final good is normalized to one. First order maximizing conditions give the demand curve and inverse demand function

$$\hat{x}_i(p) = \left(\frac{1-\alpha}{p} \right)^{\frac{1}{\alpha}} q_i L_y, \qquad \hat{p}_i(x_i) = (1-\alpha) \left(\frac{q_i L_y}{x_i} \right)^{\alpha} \tag{3}$$

Intermediate producers make profit

$$\pi_i^I = \hat{p}_i(\hat{x}_i)\hat{x}_i - c\hat{x}_i \tag{4}$$

where unit cost c is the amount of final output needed to produce a unit of the intermediate good.

Researchers

Researchers develop new designs for both entirely new products and for improvements to existing products. The researchers license these patented designs to intermediate producers. That license contract is determined by bargaining between the intermediate producers and researchers and depends on the profit stream that intermediate producer can realize using the design. We assume the licenses take the form of a fixed fee.

Consider first a new product. In this case, the researcher offers an exclusive license, giving the intermediate producer a monopoly.[5] First order maximization of (4) by

[5] It can be readily shown that this is the researchers best licensing strategy.

intermediate producers yields the monopoly price for the intermediate good as a markup over cost

$$p^M = \frac{c}{1-\alpha}, \tag{5}$$

and monopoly profits

$$\pi^N = \pi^M(q_i) = \alpha q_i L_y z, \quad z \equiv (1-\alpha)^{1/\alpha}\left(\frac{1-\alpha}{c}\right)^{\frac{1}{\alpha}-1}. \tag{6}$$

Note that the profit is proportional to the quality of the product. The researcher who designs this new product can offer it to any intermediate producer. Since the bidding for an exclusive license is competitive, we assume that the researcher extracts rent equal to the value of a stream of profits π^N.

Matters are different when a new version of product *i* is introduced that has higher quality than the version produced by the incumbent firm. Let firm 0 be the incumbent firm and let the new design have quality $q_{i1} = \Delta q_i + q_{i0} > q_{i0}$. There are three different outcomes depending on the magnitude of the quality improvement, $\Delta q_i / q_{i0}$. First, consider drastic improvements (see the Appendix for the conditions for the different outcomes). For the largest quality improvements, the entrant firm drives the incumbent out of the market and earns monopoly rents, $\pi^M(q_{i1})$. The researcher extracts these rents as above.

If, instead, the quality improvement is somewhat smaller, the innovation is non-drastic. In this case, the researcher could license the design to a firm that competes with the incumbent. Then the intermediate producer charges a limit price, p^L, earning profit of π^L (see Appendix); the incumbent earns zero and drops out of the market. But the researcher does not choose to license to a new competitor. Instead, the researcher licenses the innovation to the *incumbent*, and both parties make out better, hence this is their preferred choice. With this license, the incumbent charges the monopoly price, p^M, and earns profit $\pi^M(q_{i1})$ less the license fee. The joint profit is larger than π^L, so there is bargaining surplus, and the limit pricing profit, π^L, serves as the researcher's "threat point" in the bilateral bargaining negotiation. Assuming equal bargaining power, the researcher's profit stream is

$$\pi^R(q_{i0}, q_{i1}) = \frac{1}{2}\pi^M(q_{i1}) + \frac{1}{2}\pi^L(q_{i0}, q_{i1}). \tag{7}$$

However, this bargain is subject to a "patent troll" problem: it implies that a researcher with only a trivial improvement can lay claim to over half of the incumbent's

profits. To make the model more realistic, let incumbent firms "invent around" a troll patent at some cost for R&D and legal expenses. Let the cost of inventing around be proportional to the improvement, $F = k\Delta q_i$, where k denotes the proportional cost of inventing around and is a constant. If the incumbent can make more than π^R by inventing around, then the researcher's best plan is to offer a license fee of $V^T(\Delta q_i) = BF$, where B is the researcher's bargaining power.[6] Importantly, we don't want the cost of inventing around to be so low that incumbents can credibly threaten to invent around major improvement innovations. Evidence suggests that the patent troll model applies mainly to minor inventions and that licensees get relatively benefit from them (Lemley and Feldman 2016).[7] In the Appendix we show that there is a threshold, T, where inventing around is credible for the incumbent if $\Delta q/q_0 < T(k)$ and not for larger improvements. The threshold is decreasing in k, the cost of inventing around. Figure 1 shows the level of invent-around cost (scaled) required for a given threshold T.[8] For example, for the invent-around zone to cover 20% improvements or smaller, the scaled value of k needs to be 3 or 4.

Thus, there are four different payoffs to researchers: monopoly profits on new products, monopoly profits on drastic innovations, a "limit price" bargain, and an "invent-around" bargain. In all of these cases, the intermediate producer charges the monopoly price and produces the monopoly level of each intermediate. Inserting this monopoly level in equation (1), aggregate output becomes

$$Y = L_y^{\alpha} \int_0^N q_i^{\alpha} \hat{x} (p^M)^{1-\alpha}\, di = \frac{QL_y z}{1-\alpha}, \qquad Q \equiv \int_0^N q_i\, di\,. \tag{8}$$

Aggregate quality corresponds to aggregate productivity.

Innovation

The difference between the value of new product licenses and product improvement licenses affects the set of ideas that researchers choose to develop. We show that, compared

[6] The bargaining surplus is F and the researcher's alternative is 0. We allow the bargaining power to differ between incumbent and researcher here (and not in (7)) because we wish to explore factors that might affect it below.

[7] Many patent trolls may also benefit from improperly granted patents that are not real inventions.

[8] The y-axis is scaled by $B(r+d)/\alpha L_y z$.

to new product innovations, researchers will choose to develop relatively more low-quality improvement innovations and, on average, improvement innovations deliver less quality than do new products.

Each period, there is an exogenous, fixed flow of M technological opportunities—possible ideas that can be developed. Researchers must decide which of these technological opportunities are worthwhile expending resources to develop. To do this, they examine the quality of the resulting product. The quality of each technological opportunity, q_i for a new product or Δq_i for an improvement, is a random draw from a common distribution with a minimum support $q_i \geq q_m$. Each improvement opportunity is randomly assigned to one of the existing products with quality q_{i0}. We assume that the distribution of qualities is stationary and is the same for new products as for improvements—technologically speaking, an improvement could *be* a new product, just that it substitutes for an already existing product.

Researchers choose to develop only *feasible* opportunities, that is, opportunities where the researcher's expected profit rate is at least as large as the alternative that the researcher could earn as a production worker, namely, the wage w. We ignore "racing;" only one researcher develops each feasible idea.

The researcher's profit rate is the license fee the researcher earns times the rate at which an idea is developed. Following Romer (1990), let the rate of development equal δA, where δ represents research productivity and A is the stock of existing ideas/designs. Because ideas are non-rival, researchers can draw on existing knowledge so that a larger stock of ideas facilitates the development of new designs. Given discount rate r and obsolescence rate, d, the feasibility conditions for an innovation are then[9]

$$\frac{\pi^k(q_i)}{r+d}\delta A \geq w, \quad k = N, R, \qquad V^T(\Delta q_i)\delta A \geq w. \tag{9}$$

This feasibility condition can be used to derive the minimum quality needed for an idea to be feasible. For a new product innovation, (9) implies $q_i \alpha L_y z \frac{\delta A}{r+d} \geq w$ so that the minimally feasible quality is

[9] We assume a continuum of infinitely lived consumers, with identical intertemporally additive preferences defined over lifetime consumption, so that there is a constant interest rate, $r > 0$.

$$q_i \geq \underline{q}_N \equiv w\frac{r+d}{\alpha L_y z \delta A}. \tag{10}$$

This feasibility condition is equivalent to the "labor arbitrage" condition in other growth models, but it is only binding on the researchers developing marginally feasible opportunities; other researchers earn more than the wage.

The feasibility threshold for improvements depends on their magnitude, including whether they fall into the troll zone or the limit pricing bargain zone. We infer that because patent trolls exist, patent trolling must be feasible for some improvement innovations, so the minimum feasible improvement must generally fall into the zone where inventing around is credible. The innovation feasibility condition in this zone is $V^T(\Delta q_i)\delta A \geq w$ or[10]

$$\underline{\Delta q}_R \equiv \frac{w}{Bk\delta A} = \frac{\underline{q}_N}{K}, \quad K \equiv \frac{Bk(r+d)}{\alpha L_y z}. \tag{11}$$

Note that $K > 1$ so that $\underline{\Delta q}_R < \underline{q}_N$ as long as the invent-around zone is limited to relatively small improvements, e.g., $T < 1$. Moreover, $\underline{\Delta q}_R < \underline{q}_N$ holds even if the innovation feasibility threshold falls in the limit price zone (see Appendix). This result is not dependent on assumptions about patent trolling.

This property of , $\underline{\Delta q}_R < \underline{q}_N$ is important for our analysis. Loosely speaking, replacement innovations garner a share of the incumbent's rents—"business stealing." Since the qualities of both new product and improvement innovations are drawn from the same distribution, it means that feasible improvement innovations include relatively more lower quality innovations. And this implies that the mean improvement innovation is of lower quality than the mean new product innovation. Generally, the mean of a distribution truncated from below will be higher if the truncation cutoff is higher. To be concrete, it is helpful to assume that the q_i are drawn from a Pareto distribution with minimum value q_m and tail parameter $\beta > 1$. The cumulative distribution function for the upper tail is $\text{Prob}[q \geq q_i] = G(q_i) = \left(\frac{q_M}{q_i}\right)^{\beta}$.[11] The Pareto distribution provides the long tail of returns to R&D that is observed in the empirical literature. It also corresponds to an "innovation

[10] Remember that the researcher's payoff in an "invent-around" bargain is given by $V^T(\Delta q_i) = Bk\Delta q_i$, where B is researcher's bargaining power, and k is the proportional cost of inventing around innovation. Therefore, the feasibility condition is given by $Bk\Delta q_i\, \delta A \geq w$. Substituting w using equation (10) yields equation (11).
[11] The standard CDF for the *lower tail* Pareto distribution is $\text{Prob}[q \leq q_i] = 1 - G(q_i)$.

possibilities frontier" with a constant elasticity, as is also common in the empirical literature, and it is used as well in theoretical work (e.g., S. S. Kortum 1997). The mean of a Pareto distribution truncated below at x is

$$E[q \mid q > x] = \frac{x\beta}{\beta - 1},$$

so the mean qualities of feasible new product and replacement innovations are, respectively,

$$\bar{q}_N = \frac{\beta \underline{q}_N}{\beta - 1}, \qquad \overline{\Delta q}_R = \frac{\beta \underline{\Delta q}_R}{\beta - 1} = \frac{\bar{q}_N}{K}. \tag{12}$$

The feasibility thresholds also determine the number and mix of researchers developing technological opportunities. Given a flow of M technological opportunities, let the share of these opportunities that are from improvement innovations be ψ. For the moment we treat ψ as exogenous but will consider endogenous possibilities below. Then,

$$L_n(\underline{q}_N) = (1 - \psi)MG\left(\underline{q}_N\right) + \psi MG\left(\underline{q}_N/K\right) \tag{13}$$

where $G\left(\underline{q}_N\right)$ and $G\left(\underline{q}_N/K\right)$ is the fraction of feasible innovations among the flow of new technological opportunities and improvement technological opportunities, respectively.[12] We define the share of improvement innovations, as follows:

$$\phi = \frac{\psi MG\left(\underline{q}_N/K\right)}{\psi MG\left(\underline{q}_N/K\right) + (1 - \psi)MG\left(\underline{q}_N\right)} = \frac{1}{1 + \frac{1 - \psi}{\psi K^{\beta}}}$$

As can be seen, the replacement share of ideas that are developed, ϕ, is a function of three variables (ψ, K, β). It is increasing in the share of improvement innovation opportunities (ψ), the feasibility threshold of new innovations relative to improvement innovations (K), and the tail parameter of the Pareto distribution (β). Given ϕ, the growth of new products is defined as $\dot{N} = (1 - \phi)\dot{A}$, where the flow of all new feasible ideas is given by $\dot{A} = \delta A L_n(\underline{q}_N)$.

[12] Here, we assume that researchers dispose of non-feasible ideas and draw new ideas immediately. We do not allow researchers to dispose of feasible ideas and wait for higher-quality ideas to be drawn.

where ϕ is the replacement share of ideas that are developed. Then the growth in innovations is $\dot{A} = \delta A L_n(\underline{q}_N)$, the growth of new products is $\dot{N} = (1-\phi)\dot{A}$.

These equations allow us to calculate the obsolescence rate. The share of incumbent products that are replaced is $\frac{\phi\dot{A}}{N}$. The incumbent firm does not necessarily lose all quasi-rents when its product is obsolesced—the limit price bargain and the invent-around bargain allow the incumbent to continue to earn rents at a reduced rate. To keep the analysis simple, suppose that on average, the incumbent loses share $0 < s \leq 1$ of its rents on an obsolescence event. Then the obsolescence rate is

$$d(L_n) = s\frac{\phi\dot{A}}{N} = \frac{s\phi\delta L_n}{N}. \tag{14}$$

We close the system by setting the wage equal to the marginal productivity of production labor in equation (1),

$$w = \frac{\partial Y}{\partial L_y} = \alpha\frac{Y}{L_y} = \frac{\alpha Q z}{1-\alpha}.$$

Using (10), we get

$$\underline{q}_N = \frac{r + d(\underline{q}_N)}{(1-\alpha)\delta L_y(\underline{q}_N)} \cdot \frac{Q}{A}. \tag{15}$$

Growth

We can write the change in aggregate quality

$$\dot{Q} = \bar{q}_N\dot{N} + \overline{\Delta q}_R\left(\dot{A} - \dot{N}\right) = \bar{q}_N\dot{A}\left(1 - \phi\frac{K-1}{K}\right)$$

and the growth rate, dividing by Q and using (15),

$$g_Q = \frac{\dot{Q}}{Q} = v\frac{\dot{A}}{A}, \quad v \equiv \frac{\frac{\beta}{\beta-1}(r + d(L_n))\left(1 - \phi\frac{K-1}{K}\right)}{(1-\alpha)\delta(L - L_n)}. \tag{16}$$

In general, the growth in productivity does not equal the growth in innovations. On the balanced growth path, they are equal, however.

We impose balanced growth by assuming that aggregate intermediates, $X = \int_0^N \hat{x}(p^M)di = N\hat{x}(p^M)$ grows at a constant rate, $g_X = \dot{X}/X = \dot{N}/N = const.$ But $\dot{N} = (1-\phi)\dot{A}$, so $g_N = g_A = const.$ From this we can infer that along the balanced growth

path, $N/A = 1 - \phi$. Note also that since $g_A = \delta L_n(\underline{q}_N)$, $\underline{q}_N$ and L_n must also be constants as well as $\underline{\Delta q}_R$ and L_y along the balanced growth path. Inspecting equation (15), if $\underline{q}_N$ is constant, then Q/A must also be constant, implying that $g_Q = q_A$. Finally, since L_y is constant along the balanced growth path, equation (8) implies

$$g_Y = g_Q = g_N = g_X = g_A = \delta L_n\left(\underline{q}_N\right). \tag{17}$$

This corresponds to Romer's 1990 model. Note also that $g_Q = q_A$ implies that $v = 1$ on the balanced growth path. We can solve this at $L_n = \tilde{L}_n$ (see Appendix). A solution exists as long as L is sufficiently large.

Rising replacement

Declining growth

What happens when an exogenous increase in the replacement share, ϕ, disturbs this steady state? The steady state growth rate declines. It is straightforward to show (see Appendix), that

$$\frac{\partial \tilde{L}_n}{\partial \phi} < 0$$

for sufficiently large *L*. The intuition behind this result is that a rising replacement rate corresponds to a rising obsolescence rate (14). Higher obsolescence reduces the returns to R&D, all else equal, reducing research labor which, in turn, reduces the balanced path growth rate. There is also a modest countervailing effect: a rise in ϕ reduces the marginal productivity of production labor and the wage; this reduces L_y which correspondingly increases L_n. However, with sufficiently large L, the obsolescence effect dominates.

This analysis implies that in the long run, a rise in the replacement rate reduces the growth rate and the level of research along the balanced growth path. But the secular rise in the obsolescence rates over the last 30 years is not consistent with balanced growth (Ando, Bessen, and Wang 2025). To explore the short-term effect of rising replacement, it is helpful to see what happens to equation (16) when, starting on the balanced growth path, the replacement rate, ϕ, rises. There are two effects: 1) greater ϕ directly decreases v, and 2) greater ϕ increases the obsolescence rate, decreasing L_n. However, the latter effect is

counterfactual. As in Figure 1, the investment in R&D has not decreased. This could be because the rise in replacement has been accompanied by a rise in research productivity, δ, so that the number of researchers has not declined. To simplify the exercise, let us assume that $L_n(\underline{q}_N)$ is unchanged. To the extent that rising obsolescence does decrease the number of researchers, the drop in productivity growth will be larger.

Given this assumption, we can calculate the short-term effect of rising replacement from the change in $\left(1-\phi\frac{K-1}{K}\right)$. Assuming that the economy was on the balanced growth path in 1995, we compare $\left(1-\phi_t\frac{K-1}{K}\right)/\left(1-\phi_{1995}\frac{K-1}{K}\right)$, letting $K=3$, consistent with a modest invent-around zone (Figure 1). To estimate ϕ, we use estimated obsolescence rates (smoothed) from Ando, Bessen, and Wang (2025) and real R&D for the US private nonfarm business sector from the BEA, we calculate the net R&D capital stocks and net R&D investment. Then interpreting $\phi = 1-\frac{net\ R\&D}{gross\ R\&D}$, we get crude estimates of this ratio showing the relative reduction in the productivity growth rate (see Table 1 and Figure 2). This exercise demonstrates the rising obsolescence could have a considerable impact on productivity growth.

Moreover, this slowdown could last decades. Following a one-time jump in ϕ, the economy converges back to the balanced growth path as N/A converges to $1-\phi$. If we assume that $g_A = .013$, the BLS productivity growth rate from 1960 to 1980, then the half-life of the productivity decline (the time for the difference between N/A and $1-\phi$ to fall by half) is 53 years.[13] While this numerical exercise is hardly definitive, it suggests that rising obsolescence could be a substantial contributor to the productivity slowdown.

This analysis stands in contrast to Bloom et al. (2020) who use an equation like (17) to argue that research productivity can be measured as $\delta = g_Q/L_n$. They observe that productivity growth, g_Q, has not increased while the number of researchers, L_n, has grown substantially, so they conclude that research productivity, δ, must have declined. But equation (17) holds only along the balanced growth path. Several decades of secularly rising

[13] $N_0/A_0 = 1-\phi_0$ at $t=0$. Then let the replacement rate change to ϕ_1. Solving the differential equations for $\dot{A}$ and $\dot{N}$, dividing, and normalizing the integration constants $N_t/A_t = 1-\phi_1+(\phi_1-\phi_0)e^{-g_A t}$. Half of the transition is realized at time $t_{half} = \ln 2/g_A$. When $g_A = .013$, then the half-life is 53 years.

obsolescence rates is simply not consistent with balanced growth, hence the relevance of this equation is questionable. Our analysis of equation (16) suggests that productivity might decline with stable or rising research effort for long periods off the balanced growth path. The direction of innovation matters as well as the intensity of R&D investment.

Social welfare and policy

Just because rising replacement reduces productivity growth it does not follow that rising replacement necessarily reduces social welfare. This is because researchers' incentives to develop innovations may be less than the social return on those inventions. We can compare researchers' incentives at the feasibility margin with the marginal output of those inventions.

For new products, the researcher's incentive for a marginally feasible product is $\pi^{new} = \alpha \underline{q}_N L_y z$ at the feasibility margin; looking at equation (8), the marginal contribution to output is $\underline{q}_N L_y z$. Since $\alpha < 1$, the private incentive is less than the optimal social incentive. Some technological opportunities that would be socially beneficial are not developed. This is the frequent rationale for policies that subsidize innovative activity.

For replacement products, however, at the feasibility threshold the private present value of benefits is $V\left(\underline{\Delta q}_R\right) = Bk\underline{\Delta q}_R$. The present value of the innovation to society is $W = \frac{\underline{\Delta q}_R L_y z}{r+d}$. The ratio is $V/W = Bk(r+d)/L_y z = \alpha K$. If $\alpha \approx 0.8, K \approx 3$, both reasonable values, then the private incentive for minor improvements at the feasibility threshold exceeds the social benefit. In this case, R&D expended on minor innovations is socially wasteful and a shift to a higher rate of replacement not only lowers productivity growth but also reduces social welfare.

One policy solution to this problem is well understood in patent law: screening for patents requires that they meet a minimum "inventive step" or a non-obviousness requirement (Meurer and Strandburg 2008). Less availability of patents for trivial inventions would shift the bargaining over non-drastic innovations to a lower value of B, thus raising the minimum feasible quality of replacement innovations. Other policy changes have sought to reduce the enforceability of marginal patents.

The rise of replacement

Why has the replacement share of innovation risen? Most likely a variety of factors have contributed to the trend over the last several decades. New technology may have reduced the cost of developing minor improvements; for instance, it has been argued that software plays this role (Bessen 2022). Technology might also have become more complex, with relatively more opportunities for sequential improvement. Also, patent policy may have shifted to provide stronger rights for minor inventions, raising the business stealing effect. Patent scholars have contended that changes in the Federal courts have made it easier to obtain patents on trivial inventions and have strengthened the enforcement of these patents. (Jaffe and Lerner 2004; Bessen and Meurer 2008). Indeed, some evidence suggests that the share of patents for incremental innovations has increased relative to the share of radical or breakthrough patents (Park, Leahey, and Funk 2023; Kalyani 2022).

Another possibility is raised by Olsson (2005) who hypothesizes that waves of new, unintentionally-developed fundamental technologies give rise to a changing balance of technological opportunities between incremental and radical innovations. It is helpful to consider technological opportunities arising in a product space—a space where each location, defined by a set of product characteristics, represents a distinct product that could conceivably be developed using current technology or which has already been developed. Let the product space have Z locations. Technological opportunities arise when researchers realize possible innovations to improve one of the existing N products or to create a new product at one of the $Z - N$ undeveloped locations. If we assume that these opportunities arise randomly across the Z locations, then the share of replacement opportunities is

$$\psi = \frac{N}{Z}, \quad \phi = \frac{1}{\left(\frac{Z}{N} - 1\right) K^{-\beta} + 1}. \tag{18}$$

Then, given a fixed Z, the replacement rate rises as the number of products, N, grows. This provides a kind of "fishing out" story where the replacement share rises as the number of undeveloped new product opportunities dwindles.

However, there is no particular reason to assume that Z is constant. As new technologies emerge, new product possibilities arise, expanding the product space. CRISPR or Large Language Models create all sorts of new product possibilities that were not conceivable in any detail one hundred years ago. Economic historians have seen whole eras

of innovation arising from major advances in science and technology such as electricity or chemistry (see, for instance, Gordon 2017; Kelly et al. 2021). Some see these advances arising from new General Purpose Technologies that open up product possibilities in a wide range of new applications (Bresnahan and Trajtenberg 1995). This view of major advances suggests that the number of conceivable products might grow episodically, alternating between periods of rapid growth and stagnation. If so, the replacement rate may fall and rise over time, and, with it, the rate of productivity growth.

Conclusion

Aggregate productivity growth depends on not just the level of R&D, but also the direction of R&D between replacement and new product development. While Schumpeterian replacement unambiguously raises productivity, a rising share of replacement means relatively more low-productivity innovation, relatively more incremental innovations compared to breakthrough innovations. And this sort of shift can reduce aggregate productivity growth. One cannot infer that stagnant or declining productivity growth implies declining research productivity; it may, instead, reflect change in the kind of innovations researchers choose to develop.

Our analysis cautions against assuming that recent stagnation in productivity growth portends a *permanent* decline. Some observers see current performance as evidence of "fishing out"—perhaps productivity is stagnant because we have harvested all the "low hanging" technological opportunities. But even if we accept Gordon's assertion (2017) that information technology does not deliver the same productivity benefit of past waves of technology, that does not mean the next waves, some of which may already have begun, might not deliver more.

Generally, the model highlights the importance of increasing the range of technological opportunities for new products, increasing the size of the product space. And the way to do this likely involves greater effort into science and basic research. With a new wave of technological opportunities, the replacement share of innovation will decline, the new product share will rise, and with it, productivity growth will rise both in the long and short term. Furthermore, policies that discourage trivial innovations or limit the enforcement of patents on minor innovations may help shift the direction of research to new products and more fundamental advances.

Appendix

Bargaining solution zones for improvement innovations

First, we obtain p^L by considering prices when the final goods producer will be indifferent between the two products. From (2) and (3), the final goods producer's profit on product *i* at the profit maximizing quantity is

$$\pi_i(\hat{x}_i) = L_y^{\alpha} q_i^{\alpha} \hat{x}_i^{1-\alpha} - p_i \hat{x}_i = q_i L_y \alpha \left(\frac{1-\alpha}{p_i}\right)^{\frac{1-\alpha}{\alpha}}$$

From this we can see that the final goods producer is indifferent between two products of different quality if

$$\frac{q_0^{\frac{\alpha}{1-\alpha}}}{p_0} = \frac{q_1^{\frac{\alpha}{1-\alpha}}}{p_1}.$$

At the limit price, firm 0 can charge c or less, so the limit price is

$$p^L = c\left(\frac{q_1}{q_0}\right)^{\frac{\alpha}{1-\alpha}}.$$

If $p^L < p^M$, the incumbent can make positive profits when the entrant charges the monopoly price. This gives rise to non-drastic innovation—the entrant powers its price to p^L to drive the incumbent out of the market. Thus, non-drastic innovation occurs when

$$q_1/q_0 < (1-\alpha)^{1-\frac{1}{\alpha}}.$$

Small improvements will be non-drastic; larger ones are drastic.

If the improvement falls into the non-drastic zone, there are two bargained solutions: one where the researcher can credibly threaten a limit price alternative and one where the incumbent can credibly threaten to invent around the researcher's patent. The choice depends on the present value the incumbent expects under each alternative.

In the case of bargaining where the limit price is the alternative, following the text, the incumbent expects present value

$$W^R = \frac{1}{2(r+d)}(\pi^M - \pi^L).$$

Using the expression for the limit price above,

$$\pi^L = (p^L - c)\hat{x}(p^L) = q_0 L_y z\left(1 - \left(\frac{q_0}{q_1}\right)^{\frac{\alpha}{1-\alpha}}\right)$$

so that

$$W^R = \frac{L_y z}{2(r+d)}\left[\alpha q_1 - q_0\left(1 - \left(\frac{q_0}{q_1}\right)^{\frac{\alpha}{1-\alpha}}\right)\right].$$

In the case where inventing around is the alternative, the value to the incumbent is

$$W^T = \frac{\pi^M}{r+d} - BF = \frac{\alpha q_1 L_y z}{r+d} - Bk(q_1 - q_0).$$

The incumbent will be indifferent between the invent-around bargain and the limit price bargain when $W^T = W^R$. Figure 1 shows indifference curves for different values of $T = \Delta q/q_0$ showing $K \equiv k\frac{B(r+d)}{\alpha L_y z}$ on the y-axis. T represents the upper limit of the invent-around zone. The incumbent prefers inventing around to the left of this point and it prefers the limit price bargain to the right.

Feasibility threshold for the limit price bargaining zone

The researcher's payoff under the limit price bargain is

$$\pi^R(q_{i0}, q_{i1}) = \frac{1}{2}\pi^M(q_{i1}) + \frac{1}{2}\pi^L(q_{i0}, q_{i1}).$$

We approximate this for small improvements, $\Delta q \ll q_0$ or $q_1/q_0 \approx 1$, using the Taylor series expansion for the limit price profit,

$$\pi^L \approx L_y z\frac{\alpha}{1-\alpha}\Delta q.$$

This yields

$$\pi^R \approx \frac{1}{2}\alpha L_y z\left(q_1 + \frac{\Delta q}{1-\alpha}\right) = \frac{1}{2}\alpha L_y z\left(q_0 + \frac{\Delta q(2-\alpha)}{1-\alpha}\right)$$

$$\pi^R \approx \alpha L_y z\frac{\Delta q(2-\alpha)}{1-\alpha} + \frac{1}{2}\alpha L_y z q_0.$$

Then following the same logic as in the text, the corresponding feasibility threshold is

$$\underline{\Delta q}_{RL} \equiv w\frac{r+d}{\alpha L_y z \delta A}\frac{1-\alpha}{2-\alpha} - \frac{1}{2}\alpha L_y z q_0 < \underline{q}_N.$$

Solution for L_n

Solving $v = 1$ for L_n

$$\tilde{L}_n = \frac{(1-\phi)\big((1-\alpha)(\beta-1)\delta KL - \beta(K(1-\phi)+\phi)r\big)}{\delta\left(\beta\phi^2 s + K(1-\phi)\big((\beta-1)(1-\alpha)+\beta\phi s\big)\right)}.$$

$\tilde{L}_n > 0$ if

$$L > \frac{\beta(K(1-\phi)+\phi)r}{(1-\alpha)(\beta-1)\delta K}.$$

$$\frac{\partial \tilde{L}_n}{\partial \phi} = \frac{\left[\begin{array}{c}(1-\phi)(\beta(K(1-\phi)+\phi)r - (1-\alpha)(\beta-1)\delta KL)\cdot \\ (2\beta\phi s + K(\beta s(1-2\phi)) + \\ \big(\beta(K-1)(1-\phi)r(\beta\phi^2 s + K(1-\phi)((\beta-1)(1-\alpha)+\beta\phi s)\big) + \\ (\beta(K(1-\phi)+\phi)r - (1-\alpha)(\beta-1)\delta KL)\cdot \\ (\beta\phi^2 s + K(1-\phi)((\beta-1)(1-\alpha)+\beta\phi s)\end{array}\right]}{\delta(\beta\phi^2 s + K(1-\phi)((\beta-1)(1-\alpha)+\beta\phi s)^2}$$

Taking the limit as L grows large,

$$\lim_{L\to\infty} \frac{\partial \tilde{L}_n}{\partial \phi} = -\frac{(1-\alpha)(\beta-1)\beta KL(K(1-\phi)^2 + (2-\phi)\phi)s}{(\beta\phi^2 s + K(1-\phi)((\beta-1)(1-\alpha)+\beta\phi s)^2} < 0.$$

Hence, for large enough L, an increase in the replacement share decreases the number of researchers.

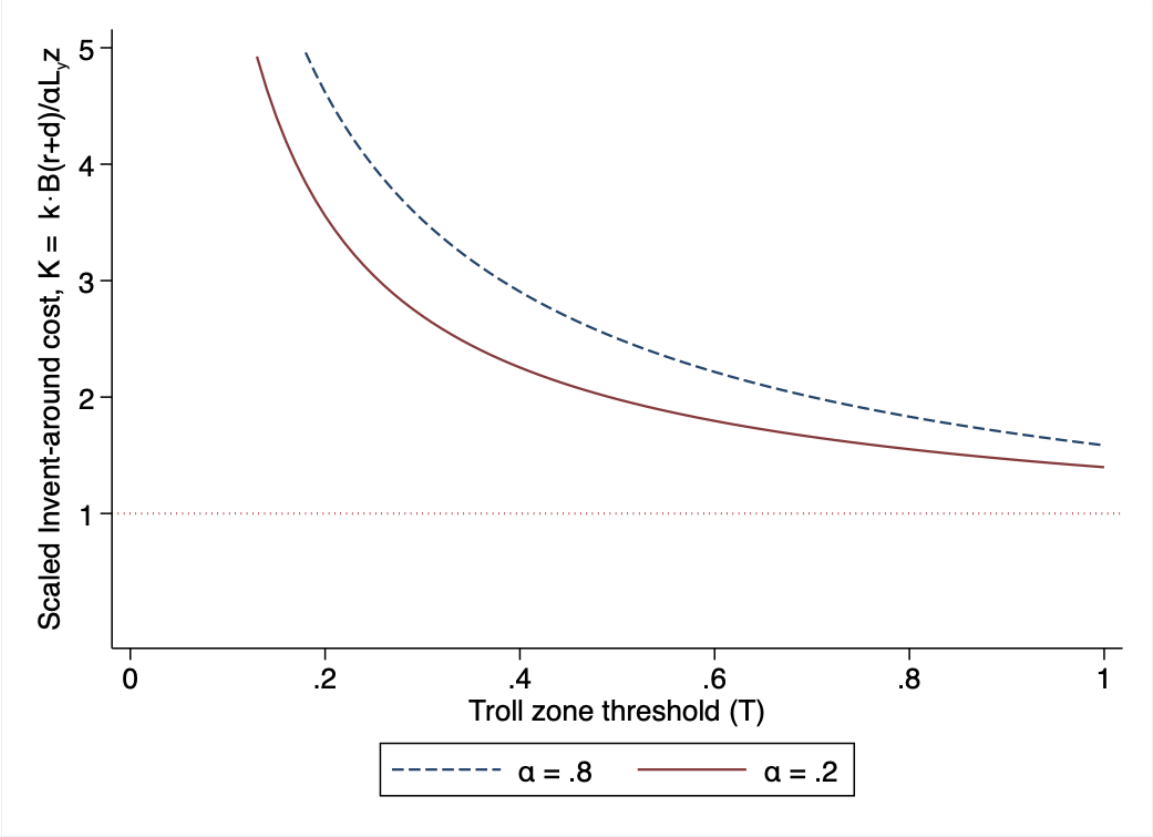
Table 1. Estimating Effect of Obsolescence on Productivity Growth

	Private fixed investment in R&D, $2017 B	Smoothed obsolescence rate	R&D Stock	net r&d	net/ gross	$\left(1-\phi_t\frac{K-1}{K}\right)/\left(1-\phi_{1995}\frac{K-1}{K}\right)$
1976	66.5	0.09	536.4	20.0	30%	0.858
1977	70.1	0.09	558.3	21.9	31%	0.871
1978	75.3	0.09	583.4	25.1	33%	0.894
1979	81.2	0.09	612.1	28.7	35%	0.916
1980	86.5	0.09	643.5	31.5	36%	0.927
1981	92.1	0.09	677.7	34.2	37%	0.935
1982	96.9	0.09	713.6	35.9	37%	0.934
1983	102.3	0.0885	752.8	39.1	38%	0.947
1984	112.8	0.0869	800.1	47.4	42%	0.988
1985	122.4	0.0854	854.2	54.1	44%	1.012
1986	127.4	0.0838	910.0	55.8	44%	1.007
1987	129.8	0.0823	964.9	54.9	42%	0.991
1988	136.9	0.0808	1023.8	58.9	43%	0.999
1989	146.2	0.0792	1088.9	65.1	45%	1.015
1990	155.5	0.0777	1159.8	70.9	46%	1.027
1991	166.7	0.0762	1238.2	78.4	47%	1.042
1992	169.9	0.0746	1315.7	77.5	46%	1.027
1993	169.2	0.0731	1388.7	73.1	43%	1.001
1994	170.2	0.0715	1459.6	70.8	42%	0.984
1995	179.6	0.07	1537.0	77.5	43%	1.000
1996	196.6	0.08	1610.7	73.7	37%	0.939
1997	210.2	0.09	1675.9	65.2	31%	0.869
1998	224.4	0.1	1732.7	56.8	25%	0.807
1999	242.2	0.11	1784.4	51.6	21%	0.764
2000	261.8	0.12	1832.1	47.7	18%	0.730
2001	267.4	0.13	1861.3	29.2	11%	0.652
2002	260.5	0.14	1861.2	-0.1	0%	0.533
2003	262.6	0.151	1842.8	-18.4	-7%	0.457
2004	264.7	0.162	1809.0	-33.8	-13%	0.395
2005	279.2	0.173	1775.2	-33.7	-12%	0.402
2006	298.4	0.184	1747.0	-28.2	-9%	0.431
2007	316.0	0.195	1722.3	-24.7	-8%	0.449
2008	325.3	0.206	1692.8	-29.5	-9%	0.435
2009	317.3	0.217	1642.8	-50.0	-16%	0.362
2010	318.5	0.228	1586.7	-56.1	-18%	0.343
2011	331.8	0.239	1539.3	-47.4	-14%	0.379
2012	334.4	0.25	1488.9	-50.4	-15%	0.370
2013	357.8	0.25	1474.5	-14.5	-4%	0.489
2014	377.0	0.25	1482.8	8.4	2%	0.557
2015	390.2	0.25	1502.4	19.5	5%	0.587
2016	424.5	0.25	1551.3	48.9	12%	0.658
2017	437.5	0.25	1601.0	49.7	11%	0.656
2018	464.3	0.25	1665.0	64.0	14%	0.682

Sources: BEA NIPA; (Ando, Bessen, and Wang 2025)

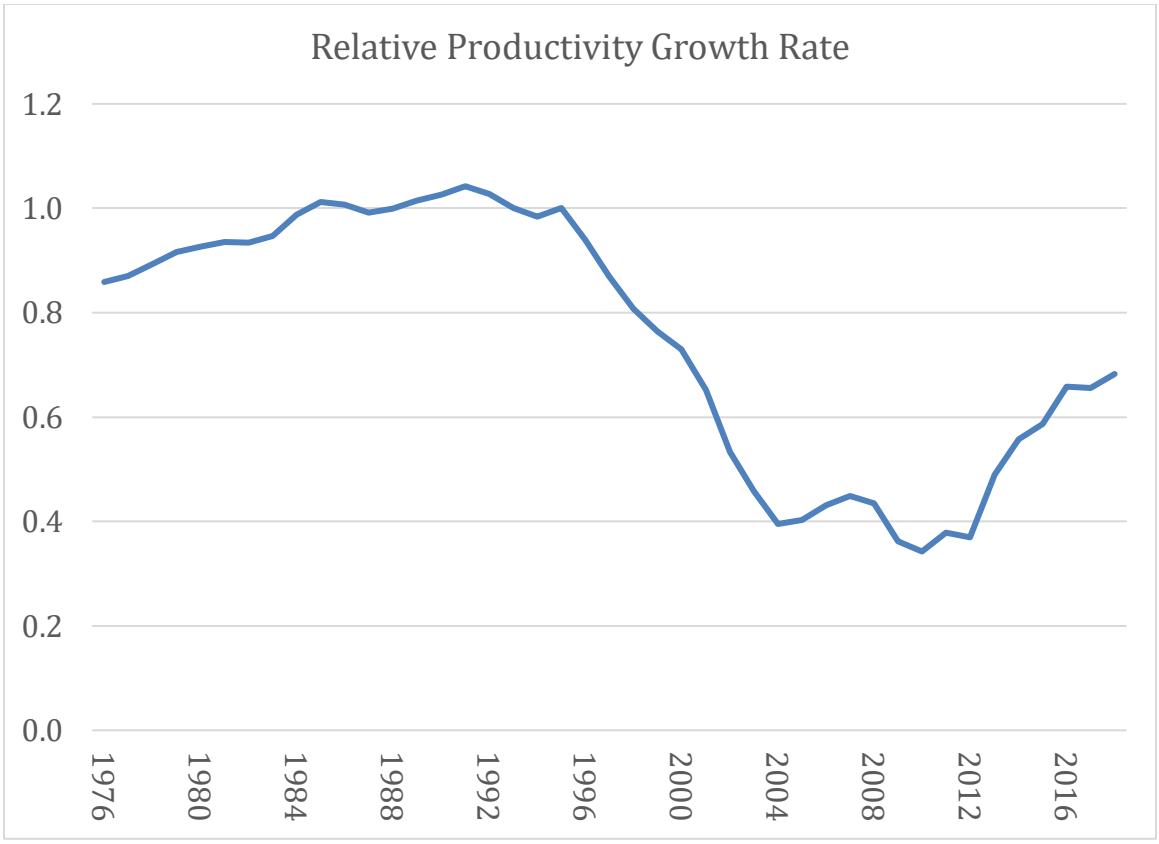
Figures

Figure 1.



Indifference curves for two values of α with $B = 0.5, r + d = 0.2, L_y z = 1$ calculated for $W^T = W^R$, using the expressions in the Appendix.

Figure 2. Estimated decrease in productivity growth (see text and Table 1)



References

Acemoglu, Daron, Ufuk Akcigit, Harun Alp, Nicholas Bloom, and William Kerr. 2018. "Innovation, Reallocation, and Growth." *American Economic Review* 108 (11): 3450–91.

Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J Klenow, and Huiyu Li. 2023. "A Theory of Falling Growth and Rising Rents." *Review of Economic Studies* 90 (6): 2675–2702.

Aghion, Philippe, and Peter Howitt. 1992. "A Model of Growth Through Creative Destruction." *Econometrica: Journal of the Econometric Society*, 323–51.

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

Ando, Yoshiki, James Bessen, and Xiupeng Wang. 2025. "The Rising Returns to R&D."

Arrow, Kenneth. 1962. "Economic Welfare and the Allocation of Resources for Invention." In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609–26. Princeton University Press.

Bessen, James. 2022. *The New Goliaths: How Corporations Use Software to Dominate Industries, Kill Innovation, and Undermine Regulation*. New Haven: Yale University Press.

Bessen, James, and Michael James Meurer. 2008. *Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk*. Princeton University Press.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review* 110 (4): 1104–44.

Boeing, Philipp, and Paul Hünermund. 2020. "A Global Decline in Research Productivity? Evidence from China and Germany." *Economics Letters* 197:109646.

Bresnahan, Timothy F, and Manuel Trajtenberg. 1995. "General Purpose Technologies 'Engines of Growth'?" *Journal of Econometrics* 65 (1): 83–108.

Caballero, Ricardo J, and Adam B Jaffe. 1993. "How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth." *NBER Macroeconomics Annual* 8:15–74.

Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J Klenow. 2019. "How Destructive Is Innovation?" *Econometrica* 87 (5): 1507–41.

Gordon, Robert J. 2017. *The Rise and Fall of American Growth: The US Standard of Living since the Civil War*. Princeton University Press.

Griliches, Zvi. 2009. "Productivity, R&D, and the Data Constraint." In *The Economic Impact of Knowledge*, 213–41. Routledge.

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

Howitt, Peter. 1999. "Steady Endogenous Growth with Population and R. & D. Inputs Growing." *Journal of Political Economy* 107 (4): 715–30.

Jaffe, Adam B, and Josh Lerner. 2004. *Innovation and Its Discontents: How Our Broken Patent System Is Endangering Innovation and Progress, and What to Do about It*. Princeton University Press.

Jones, Charles I, and John C Williams. 2000. "Too Much of a Good Thing? The Economics of Investment in R&D." *Journal of Economic Growth* 5:65–85.

Kalyani, Aakash. 2022. "The Creativity Decline: Evidence from US Patents." *Available at SSRN 4318158*.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. "Measuring Technological Innovation over the Long Run." *American Economic Review: Insights* 3 (3): 303–20.

Kortum, Samuel. 1993. "Equilibrium R&d and the Patent–R&d Ratio: Us Evidence." *The American Economic Review* 83 (2): 450–57.
Kortum, Samuel S. 1997. "Research, Patenting, and Technological Change." *Econometrica: Journal of the Econometric Society*, 1389–1419.
Lemley, Mark A, and Robin Feldman. 2016. "Patent Licensing, Technology Transfer, and Innovation." *American Economic Review* 106 (5): 188–92.
Li, Chol-Won. 2000. "Endogenous vs. Semi-Endogenous Growth in a Two-R&D-Sector Model." *The Economic Journal* 110 (462): 109–22.
Liu, Ernest, Atif Mian, and Amir Sufi. 2022. "Low Interest Rates, Market Power, and Productivity Growth." *Econometrica* 90 (1): 193–221.
March, James G. 1991. "Exploration and Exploitation in Organizational Learning." *Organization Science* 2 (1): 71–87.
Meurer, Michael J, and Katherine J Strandburg. 2008. "Patent Carrots and Sticks: A Model of Nonobviousness." *Lewis & Clark L. Rev.* 12:547.
Mokyr, Joel. 1992. *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press.
Olsson, Ola. 2005. "Technological Opportunity and Growth." *Journal of Economic Growth* 10:31–53.
Park, Michael, Erin Leahey, and Russell J Funk. 2023. "Papers and Patents Are Becoming Less Disruptive over Time." *Nature* 613 (7942): 138–44.
Ridder, Maarten De. 2024. "Market Power and Innovation in the Intangible Economy." *American Economic Review* 114 (1): 199–251.
Romer, Paul M. 1990. "Endogenous Technological Change." *Journal of Political Economy* 98 (5, Part 2): S71–102.
Segerstrom, Paul S. 1998. "Endogenous Growth without Scale Effects." *American Economic Review*, 1290–1310.