

Are Ideas Getting Harder to Find?[†]

By NICHOLAS BLOOM, CHARLES I. JONES, JOHN VAN REENEN,
AND MICHAEL WEBB*

Long-run growth in many models is the product of two terms: the effective number of researchers and their research productivity. We present evidence from various industries, products, and firms showing that research effort is rising substantially while research productivity is declining sharply. A good example is Moore’s Law. The number of researchers required today to achieve the famous doubling of computer chip density is more than 18 times larger than the number required in the early 1970s. More generally, everywhere we look we find that ideas, and the exponential growth they imply, are getting harder to find. (JEL D24, E23, O31, O47)

This paper applies the growth accounting of Solow (1957) to the production function for new ideas. The basic insight can be explained with a simple equation, highlighting a stylized view of economic growth that emerges from idea-based growth models:

$$\begin{array}{rcc} \text{Economic growth} & = & \text{Research productivity} \times \text{Number of researchers.} \\ \text{e.g., 2\% or 5\%} & & \downarrow(\text{falling}) \qquad \uparrow(\text{rising}) \end{array}$$

Economic growth arises from people creating ideas. As a matter of accounting, we can decompose the long-run growth rate into the product of two terms: the effective number of researchers and their research productivity. We present a wide range of empirical evidence showing that in many contexts and at various levels of disaggregation, research effort is rising substantially, while research productivity is

* Bloom: Department of Economics, Stanford University, and NBER (email: nbloom@stanford.edu); Jones: Graduate School of Business, Stanford University, and NBER (email: chad.jones@stanford.edu); Van Reenen: Department of Economics, MIT, LSE, and NBER (email: vanreene@mit.edu); Webb: Department of Economics, Stanford University (email: mww@stanford.edu). Emi Nakamura was the coeditor for this article. We are grateful to three anonymous referees for detailed comments. We also thank Daron Acemoglu, Philippe Aghion, Ufuk Akcigit, Michele Boldrin, Ben Jones, Pete Klenow, Sam Kortum, Peter Kruse-Andersen, Rachel Ngai, Pietro Peretto, John Seater, Chris Tonetti, and seminar participants at Bocconi, the CEPR Macroeconomics and Growth Conference, CREI, George Mason, Harvard, LSE, MIT, Minneapolis Fed, NBER growth meeting, the NBER Macro across Time and Space conference, the Rimini Conference on Economics and Finance, and Stanford for helpful comments. We are grateful to Antoine Dechezlepretre, Keith Fuglie, Dietmar Harhoff, Wallace Huffman, Brian Lucking, Unni Pillai, and Greg Traxler for extensive assistance with data. The Alfred Sloan Foundation, SRF, ERC, and ESRC have provided financial support. Any opinions and conclusions expressed herein are those of the authors and do not necessarily represent the views of the US Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed. Data and replication files can be found at <http://doi.org/10.3886/E111743V1>.

[†]Go to <https://doi.org/10.1257/aer.20180338> to visit the article page for additional materials and author disclosure statements.

declining sharply. Steady growth, when it occurs, results from the offsetting of these two trends.

Perhaps the best example of this finding comes from Moore's Law, one of the key drivers of economic growth in recent decades. This "law" refers to the empirical regularity that the number of transistors packed onto a computer chip doubles approximately every two years. Such doubling corresponds to a constant exponential growth rate of 35 percent per year, a rate that has been remarkably steady for nearly half a century. As we show in detail below, this growth has been achieved by engaging an ever-growing number of researchers to push Moore's Law forward. In particular, the number of researchers required to double chip density today is more than 18 times larger than the number required in the early 1970s. At least as far as semiconductors are concerned, ideas are getting harder to find. Research productivity in this case is declining sharply, at a rate of 7 percent per year.

We document qualitatively similar results throughout the US economy, providing detailed microeconomic evidence on idea production functions. In addition to Moore's Law, our case studies include agricultural productivity (corn, soybeans, cotton, and wheat) and medical innovations. Research productivity for seed yields declines at about 5 percent per year. We find a similar rate of decline when studying the mortality improvements associated with cancer and heart disease. Finally, we examine two sources of firm-level panel data, Compustat and the US Census of Manufacturing. While the data quality from these samples is coarser than our case studies, the case studies suffer from possibly not being representative. We find substantial heterogeneity across firms, but research productivity declines at a rate of around 10 percent per year in Compustat and 8 percent per year in the Census.

Perhaps research productivity is declining sharply within particular cases and yet not declining for the economy as a whole. While existing varieties run into diminishing returns, perhaps new varieties are always being invented to stave this off. We consider this possibility by taking it to the extreme. Suppose each variety has a productivity that cannot be improved at all, and instead aggregate growth proceeds entirely by inventing new varieties. To examine this case, we consider research productivity for the economy as a whole. We once again find that it is declining sharply: aggregate growth rates are relatively stable over time,¹ while the number of researchers has risen enormously. In fact, this is simply another way of looking at the original point of Jones (1995), and we present this application first to illustrate our methodology. We find that research productivity for the aggregate US economy has declined by a factor of 41 since the 1930s, an average decrease of more than 5 percent per year.

This is a good place to explain why we think looking at the macrodata is insufficient and why studying the idea production function at the micro level is crucial. Section II discusses this issue in more detail. The overwhelming majority of papers on economic growth published in the past decade are based on models in

¹There is a debate over whether the slower rates of growth over the last decade are a temporary phenomenon due to the global financial crisis or a sign of slowing technological progress. Gordon (2016) argues that the strong US productivity growth between 1996 and 2004 was a temporary blip and that productivity growth will, at best, return to the lower growth rates of 1973–1996. Although we do not need to take a stance on this, note that if frontier TFP growth really has slowed down, this only strengthens our argument.

which research productivity is constant.² An important justification for assuming constant research productivity is an observation first made in the late 1990s by a series of papers written in response to the aggregate evidence.³ These papers highlighted that composition effects could render the aggregate evidence misleading: perhaps research productivity at the micro level is actually stable. The rise in aggregate research could apply to an extensive margin, generating an increase in product variety, so that the number of researchers per variety, and thus micro-level research productivity and growth rates themselves, are constant. The aggregate evidence, then, may tell us nothing about research productivity at the micro level. Hence, the contribution of this paper: study the idea production function at the micro level to see directly what is happening to research productivity there.

Not only is this question interesting in its own right, but it is also informative about the kind of models that we use to study economic growth. Despite large declines in research productivity at the micro level, relatively stable exponential growth is common in the cases we study (and in the aggregate US economy). How is this possible? Looking back at the equation that began the introduction, declines in research productivity must be offset by increased research effort, and this is indeed what we find.

Putting these points together, we see our paper as making three related contributions. First, it looks at many layers of evidence simultaneously. Second, the paper uses a conceptually consistent accounting approach across these layers, one derived from core models in the growth literature. Finally, the paper's evidence is informative about the kind of models that we use to study economic growth.

Our selection of cases is driven primarily by the requirement that we are able to obtain data on both the "idea output" and the corresponding "research input." We looked into a large number of possible cases to study, only a few of which have made it into this paper; indeed, we wanted to report as many cases as possible. For example, we also considered the internal combustion engine, the speed of air travel, the efficiency of solar panels, the Nordhaus (1997) "price of light" evidence, and the sequencing of the human genome. We would have loved to report results for these cases. In each of them, it was relatively easy to get an "idea output" measure. However, it proved impossible to get a series for the research input that we felt corresponded to the idea output. For example, the Nordhaus price of light series would make a great additional case. But many different types of research contribute to the falling price of light, including the development of electric generators, the discovery of compact fluorescent bulbs, and the discovery of LEDs. We simply did not know how to construct a research series that would capture all the relevant R&D. The same problem applies to the other cases we considered but could not complete. For example, it is possible to get R&D spending by the government and by a few select companies on sequencing the human genome. But it turns out that Moore's Law is itself an important contributor to the fall in the price of gene sequencing. How should we combine these research inputs? In the end, we report the cases in which

²Examples are cited after equation (1).

³The initial papers included Dinopoulos and Thompson (1998), Peretto (1998), Young (1998), and Howitt (1999); Section II contains additional references.

we felt most confident. We hope our paper will stimulate further research into other case studies of changing research productivity.

The remainder of the paper is organized as follows. After a literature review in the next subsection, Section I lays out our conceptual framework and presents the aggregate evidence on research productivity to illustrate our methodology. Section II places this framework in the context of growth theory and suggests that applying the framework to microdata is crucial for understanding the nature of economic growth. Sections III through VI consider our applications to Moore's Law, agricultural yields, medical technologies, and firm-level data. Section VII then revisits the implications of our findings for growth theory, and Section VIII concludes.

Relationship to the Existing Literature

Other papers also provide evidence suggesting that ideas may be getting harder to find. A large literature documents that the flow of new ideas per research dollar is declining. For example, Griliches (1994) provides a summary of the earlier literature exploring the decline in patents per dollar of research spending; Kogan et al. (2017) has more recent evidence; and Kortum (1993) provides detailed evidence on this point. Scannell et al. (2012) and Pammolli, Magazzini, and Riccaboni (2011) point to a well-known decline in pharmaceutical innovation per dollar of pharmaceutical research. Absent theory, these seem like natural measures of research productivity. However, as explained in detail below, it turns out that essentially all the idea-driven growth models in the literature predict that ideas per (real) research dollar will be declining. In other words, these natural measures are not really informative about whether research faces constant or diminishing returns. Instead, the right measure according to theory is the flow of ideas divided by the number of researchers (perhaps including a quality adjustment). Our paper tries to make this clear and to focus on the measures of research productivity that are suggested by theory as being most informative.

Second, many earlier studies use patents as an indicator of ideas. For example, Griliches (1994) and Kortum (1997) emphasize that patents per researcher declined sharply between 1920 and 1990.⁴ The problem with this stylized fact is that it is no longer true! For example, see Kortum and Lerner (1998) and Webb et al. (2018). Starting in the 1980s, patent grants by the USPTO began growing much faster than before, leading patents per capita and patents per researcher to stabilize and even increase. The patent literature is very rich and has interpreted this fact in different ways. It could suggest, for example, that ideas are no longer getting harder to find. Alternatively, maybe a patent from 50 years ago and a patent today mean different things because of changes in what can be patented (algorithms, software) and changes in the legal setting; see Gallini (2002), Henry and Turner (2006), and Jaffe and Lerner (2006). In other words, the relationship between patents and "ideas" may itself not be stable over time, making this evidence hard to interpret, a point made by Lanjouw and Schankerman (2004). Our paper focuses on nonpatent measures

⁴See also Evenson (1984, 1991, 1993) and Lanjouw and Schankerman (2004).

of ideas and provides new evidence that we hope can help resolve some of these questions.

Gordon (2016) reports extensive new historical evidence from throughout the nineteenth and twentieth centuries to suggest that ideas are getting harder to find. Cowen (2011) synthesizes earlier work to explicitly make the case. Benjamin Jones (2009, 2010) documents a rise in the age at which inventors first patent and a general increase in the size of research teams, arguing that over time more and more learning is required just to get to the point where researchers are capable of pushing the frontier forward. We see our evidence as complementary to these earlier studies but more focused on drawing out the tight connections to growth theory.

Finally, there is a huge and rich literature linking firm performance (such as productivity) to R&D inputs (see Hall, Mairesse, and Mohnen 2010 for a survey). Three findings from this literature are that (i) firm productivity is positively related to its own R&D, (ii) there are significant spillovers of R&D between firms, and (iii) these relationships are at least partially causal. Our paper is consistent with these three findings and our firm-level analysis in Section VI is closely tied to this body of work. We go beyond this literature by using growth theory to motivate the specific micro facts that we document and discuss these links in more detail in Section VII.

I. Research Productivity and Aggregate Evidence

A. The Conceptual Framework

An equation at the heart of many growth models is an idea production function taking a particular form:

$$(1) \quad \frac{\dot{A}_t}{A_t} = \alpha S_t.$$

Classic examples include Romer (1990) and Aghion and Howitt (1992), but many recent papers follow this approach, including Aghion, Akcigit, and Howitt (2014); Acemoglu and Restrepo (2016); Akcigit, Celik, and Greenwood (2016); and Jones and Kim (2018). In the equation above, \dot{A}_t/A_t is total factor productivity growth in the economy. The variable S_t (think “scientists”) is some measure of research input, such as the number of researchers. This equation then says that the growth rate of the economy, through the production of new ideas, is proportional to the number of researchers.

Relating \dot{A}_t/A_t to ideas runs into the familiar problem that ideas are hard to measure. Even as simple a question as “What are the units of ideas?” is troublesome. We follow much of the literature, including Aghion and Howitt (1992), Grossman and Helpman (1991), and Kortum (1997), and define ideas to be in units so that a constant flow of new ideas leads to constant exponential growth in A . For example, each new idea raises incomes by a constant percentage, on average, rather than by a certain number of dollars. This is the standard approach in the quality ladder literature on growth: ideas are proportional improvements in productivity. The patent statistics for most of the twentieth century are consistent with this view; indeed, this was a key piece of evidence motivating Kortum (1997). This definition means that the left-hand side of equation (1) corresponds to the flow of new ideas. However, this

is clearly just a convenient definition, and in some ways a more accurate title for this paper would be “Is Exponential Growth Getting Harder to Achieve?”

We can now define the productivity of the idea production function as the ratio of the output of ideas to the inputs used to make them:

$$(2) \quad \text{Research productivity} := \frac{\dot{A}_t/A_t}{S_t} = \frac{\text{number of new ideas}}{\text{number of researchers}}.$$

The null hypothesis tested in this paper comes from the relationship assumed in (1). Substituting this equation into the definition of research productivity, we see that (1) implies that research productivity equals α , that is, research productivity is constant over time. This is the standard hypothesis in much of the growth literature. Under this null, a constant number of researchers can generate constant exponential growth.

The reason this is such a common assumption is also easy to see in equation (1). With constant research productivity, a research subsidy that increases the number of researchers permanently will permanently raise the growth rate of the economy. In other words “constant research productivity” and the fact that sustained research subsidies produce “permanent growth effects” are equivalent statements.⁵ This clarifies a claim in the introduction: testing the null hypothesis of constant research productivity is interesting in its own right but also because it is informative about the kind of models that we use to study economic growth.

B. Aggregate Evidence

The bulk of the evidence presented in this paper concerns the extent to which a constant level of research effort can generate constant exponential growth within a relatively narrow category, such as a firm or a seed type or Moore’s Law or a health condition. We provide consistent evidence that the historical answer to this question is no: research productivity is declining at a substantial rate in virtually every place we look.

This finding raises a natural question, however. What if there is sharply declining research productivity *within* each product line, but growth is sustained by the creation of new product lines? First there was steam power, then electric power, then the internal combustion engine, then the semiconductor, then gene editing, and so on. Maybe there is limited opportunity within each area for productivity improvement and long-run growth occurs through the invention of entirely new areas. An analysis focused on microeconomic case studies might never reveal this to be the case.

The answer to this concern turns out to be straightforward and is an excellent place to begin. First, consider the extreme case where there is no possibility at all for productivity improvement in a product line and all productivity growth comes from adding new product lines. Of course, this is just the original Romer (1990)

⁵The careful reader may wonder about this statement in richer models: for example, lab equipment models where research is measured in goods rather than in bodies or models with both horizontal and vertical dimensions to growth. These extensions will be incorporated below in such a way as to maintain the equivalence between “constant research productivity” and “permanent growth effects.”

model itself, and to generate constant research productivity in that case requires the equation with which we started the paper:

$$(3) \quad \frac{\dot{A}_t}{A_t} = \alpha S_t.$$

In this interpretation, A_t represents the number of product varieties and S_t is the aggregate number of researchers. Even with no ability to improve productivity within each variety, a constant number of researchers can sustain exponential growth if the variety-discovery function exhibits constant research productivity.

This hypothesis, however, runs into an important well-known problem noted by Jones (1995). For the US economy as a whole, exponential growth rates in GDP per person since 1870 or in total factor productivity since the 1930s, which are related to the left side of equation (3), are relatively stable or even declining. But measures of research effort, the right side of the equation, have grown tremendously. When applied to the aggregate data, our approach of looking at research productivity is just another way of making this same point.

To illustrate the approach, we use the decadal averages of TFP growth to measure the “output” of the idea production function. For the input, we use the NIPA measure of investment in “intellectual property products,” a number that is primarily made up of research and development spending but also includes expenditures on creating other nonrival goods like computer software, music, books, and movies. As explained further below, we deflate this input by a measure of the average annual earnings for men with four or more years of college so that it measures the “effective” number of researchers that the economy’s R&D spending could purchase. These basic data are shown in Figure 1. Note that we use the same scale on the two vertical axes to reflect the null hypothesis that TFP growth and effective research should behave similarly. But of course the two series look very different.

Figure 2 shows research productivity and research effort by decade. Since the 1930s, research effort has risen by a factor of 23, an average growth rate of 4.3 percent per year. Research productivity has fallen by an even larger amount, by a factor of 41 (or at an average growth rate of -5.1 percent per year). By construction, a factor of 23 of this decline is due to the rise in research effort and so less than a factor of 2 is due to the well-known decline in TFP growth.

This aggregate evidence could be improved on in many ways. One might question the TFP growth numbers: how much of TFP growth is due to innovation versus reallocation or declines in misallocation? One might seek to include international research in the input measure.⁶ But reasonable variations along these lines would not change the basic point: a model in which economic growth arises from the discovery of newer and better varieties with limited possibilities for productivity growth *within* each variety exhibits sharply-declining research productivity. If one wishes to maintain the hypothesis of constant research productivity, one must look

⁶Online Appendix Figure 1 reports alternative R&D measures using full-time equivalent researchers rather than deflated spending. It also looks at R&D measures that include the whole OECD (rather than just the United States) and also Russia and China. Although the exact numbers change (our baseline is in the middle of the pack), there has been a substantial increase in the volume of R&D no matter which series we use.




FIGURE 1. AGGREGATE DATA ON GROWTH AND RESEARCH EFFORT

Notes: The idea output measure is TFP growth, by decade (and for 2000–2014 for the latest observation). For the years since 1950, this measure is the Bureau of Labor Statistics (2017) Private Business Sector multifactor productivity growth series, adding back in the contributions from R&D and IPP. For the 1930s and 1940s, we use the measure from Gordon (2016). The idea input measure, *Effective number of researchers*, is gross domestic investment in intellectual property products from the National Income and Product Accounts (Bureau of Economic Analysis 2017), deflated by a measure of the nominal wage for high-skilled workers.




FIGURE 2. AGGREGATE EVIDENCE ON RESEARCH PRODUCTIVITY

Notes: Research productivity is the ratio of idea output, measured as TFP growth, to the effective number of researchers. See Notes to Figure 1 and the online Appendix. Both research productivity and research effort are normalized to the value of 1 in the 1930s.

elsewhere. It is for this reason that the literature, and this paper, turns to the micro side of economic growth.

II. Refining the Conceptual Framework

In this section, we further develop the conceptual framework. First, we explain why the aggregate evidence just presented can be misleading, motivating our focus on microdata. Second, we consider the measurement of research productivity when

the input to research is R&D expenditures (i.e., “goods”) rather than just bodies or researchers (i.e., “time”). Finally, we discuss various extensions.

A. *The Importance of Microdata*

The null hypothesis that research productivity is constant over time is attractive conceptually in that it leads to models in which changes in policies related to research can permanently affect the growth rate of the economy. Several papers, then, have proposed alternative models in which the calculations using aggregate data can be misleading about research productivity. The insight of Dinopoulos and Thompson (1998), Peretto (1998), Young (1998), and Howitt (1999) is that the aggregate evidence may be masking important heterogeneity, and that research productivity may nevertheless be constant for a significant portion of the economy. Perhaps the idea production function for individual products shows constant research productivity. The aggregate numbers may simply capture the fact that every time the economy gets larger we add more products.⁷

To see the essence of the argument, suppose that the economy produces N_t different products, and each of these products is associated with some quality level A_{it} . Innovation can lead the quality of each product to rise over time according to an idea production function,

$$(4) \quad \frac{\dot{A}_{it}}{A_{it}} = \alpha S_{it}.$$

Here, S_{it} is the number of scientists devoted to improving the quality of good i , and in a symmetric case, we might have $S_{it} = S_t/N_t$. The key is that the aggregate number of scientists S_t can be growing, but perhaps the number of scientists per product S_t/N_t is not growing. This can occur in equilibrium if the number of products itself grows endogenously at the right rate. In this case, the aggregate evidence discussed earlier would not tell us anything about the idea production functions associated with the quality improvements of each variety. Instead, aggregation masks the true constancy of research productivity at the micro level.

This insight provides one of the key motivations for the present paper: to study the idea production function at the micro level. That is, we study equation (4) directly and consider research productivity for individual products:

$$(5) \quad \text{Research productivity} := \frac{\dot{A}_{it}/A_{it}}{S_{it}}.$$

B. “Lab Equipment” Specifications

In many applications, the input that we measure is R&D expenditures rather than the number of researchers. In fact, one could make the case that this is a more desirable measure in that it weights the various research inputs according to their relative prices: if expanding research involves employing people of lower talent, this will be properly measured by R&D spending. When the only input into

⁷This line of research has been further explored by Aghion and Howitt (1998), Li (2000), Laincz and Peretto (2006), Ha and Howitt (2007), Kruse-Andersen (2016), and Peretto (2016a, b).

ideas is researchers, deflating R&D expenditures by an average wage will recover a quality-adjusted quantity of researchers. In practice, R&D expenditures also include spending on capital goods and materials. As explained next, deflating by the nominal wage to get an “effective number of researchers” that this research spending could hypothetically purchase remains a good way to proceed.

In the growth literature, these specifications are called “lab equipment” models, because implicitly both capital and labor are used as inputs to produce ideas. In lab equipment models, the endogenous growth case occurs when the idea production function takes the form

$$(6) \quad \dot{A}_t = \alpha R_t,$$

where R_t is measured in units of a final output good. For the moment, we discuss this issue in the context of a single-good economy; in the next section, we explain how the analysis extends to the case of multiple products.

To see why equation (6) delivers endogenous growth, it is necessary to specify the economic environment more fully. First, suppose there is a final output good that is produced with a standard Cobb-Douglas production function:

$$(7) \quad Y_t = K_t^\theta (A_t L)^{1-\theta},$$

where we assume labor is fixed, for simplicity. Next, the resource constraint for this economy is

$$(8) \quad Y_t = C_t + I_t + R_t.$$

That is, final output is used for consumption, investment in physical capital, or research.

We can now combine these three equations to get the endogenous growth result. Dividing both sides of the production function for final output by Y_t^θ and rearranging yields

$$(9) \quad Y_t = \left(\frac{K_t}{Y_t} \right)^{\frac{\theta}{1-\theta}} A_t L.$$

Then, letting $s_t := R_t/Y_t$ denote the share of the final good spent on research, the idea production function in (6) can be expressed as

$$(10) \quad \dot{A}_t = \alpha R_t = \alpha s_t Y_t = \alpha s_t \left(\frac{K_t}{Y_t} \right)^{\frac{\theta}{1-\theta}} A_t L.$$

And rearranging gives

$$(11) \quad \frac{\dot{A}_t}{A_t} = \alpha \underbrace{\left(\frac{K_t}{Y_t} \right)^{\frac{\theta}{1-\theta}}}_{\text{research productivity}} \times \underbrace{s_t L}_{\text{“scientists”}}.$$

It is now easy to see how this setup generates endogenous growth. Along a balanced growth path, the capital-output ratio K/Y will be constant, as will the research investment share s_t . If we assume there is no population growth, then

equation (11) delivers a constant growth rate of total factor productivity in the long run. Moreover, a permanent increase in the R&D share s will permanently raise the growth rate of the economy.

Looking back at the idea production function in (6), the question is then how to define research productivity there. The answer is both intuitive and simple: we deflate the R&D expenditures R_t by the wage to get a measure of “effective scientists.” Letting $w_t = \bar{\theta}Y_t/L_t$ be the wage for labor in this economy,⁸ (6) can be written as

$$(12) \quad \frac{\dot{A}_t}{A_t} = \frac{\alpha w_t}{A_t} \times \frac{R_t}{w_t}.$$

Importantly, the two terms on the right-hand side of this equation will be constant along a balanced growth path (BGP) in a standard endogenous growth model. It is easy to see that $\tilde{S}_t := R_t/w_t = (R_t/Y_t) \cdot (Y_t/w_t) = s_t L/\bar{\theta}$. And of course w_t/A_t is also constant along a BGP.

In other words, if we deflate R&D spending by the economy’s wage rate, we get \tilde{S}_t , a measure of the number of researchers the R&D spending could purchase. Research labs spend on other things as well, like lab equipment and materials, but the theory makes clear that \tilde{S}_t is a useful measure for constructing research productivity. Hence, we will refer to \tilde{S}_t as “effective scientists” or “research effort.”

The idea production function in (12) can then be written as

$$(13) \quad \frac{\dot{A}_t}{A_t} = \tilde{\alpha}_t \tilde{S}_t$$

where both $\tilde{\alpha}_t$ and \tilde{S}_t will be constant in the long run under the null hypothesis of endogenous growth. We can therefore define research productivity in the lab equipment setup in a way that parallels our earlier treatment:

$$(14) \quad \text{Research productivity} := \frac{\dot{A}_t/A_t}{\tilde{S}_t}.$$

The only difference is that we deflate R&D expenditures by a measure of the nominal wage to get \tilde{S} . An easy intuition for (14) is this: endogenous growth requires that a constant population, or a constant number of researchers, be able to generate constant exponential growth. Deflating R&D spending by the wage puts the R&D input in units of “people” so that constant research productivity is equivalent to the null hypothesis of endogenous growth.

Equation (12) also makes clear why deflating R&D spending by the wage is important. If we did not and instead naively computed research productivity by dividing \dot{A}_t/A_t by R_t , we would find that research productivity would be falling because of the rise in A_t , even in the endogenous growth case. In other words, virtually all idea-driven growth models in the literature predict that ideas per research dollar is declining; theory suggests that ideas per researcher is a much more informative measure.

⁸The bar over θ allows for the possibility, common in these models, that labor is paid proportionately less than its marginal product because of imperfect competition.

As a measure of the nominal wage in our empirical applications, we use mean personal income from the Current Population Survey for males with a bachelor's degree or more of education.⁹

C. *Heterogeneous Goods and the Lab Equipment Specification*

In the previous two subsections, we discussed (i) what happens if research productivity is only constant within each product while the number of products grows and (ii) how to define research productivity when the input to research is measured in goods rather than bodies. Here, we explain how to put these two together.

Among the very first models that used both horizontal and vertical research to neutralize scale effects, only Howitt (1999) used the lab-equipment approach. In that paper, it turns out that the method we have just discussed, deflating R&D expenditures by the economy's average wage, works precisely as explained above. That is, research productivity for each product should be constant if one divides the growth rate by the effective number of researchers working to improve that product.¹⁰ This is true more generally in these horizontal/vertical models of growth whenever product variety grows at the same rate as the economy's population. Peretto (2016a) cites a large literature suggesting that this is the case: product variety and population scale together over time and across countries.¹¹ The two previous subsections then merge together very naturally.

D. *Diminishing Returns at a Point in Time*

One other potential modification to the idea production function that has been considered in the literature is a duplication externality. Specifically, perhaps the idea production function depends on S_t^λ , where λ is less than 1. Doubling the number of researchers may less than double the production of new ideas because of duplication or because of some other source of diminishing returns.

We could incorporate this effect into our analysis explicitly but we choose instead to focus on the benchmark case of $\lambda = 1$ for several reasons. First, our measurement

⁹See the online Appendix for more details. A shortcoming of using the college earnings series is that the increase in college participation may mean that less talented people are attending college over time. To the extent that this is true, our deflator may understate the rise in the wage for a constant-quality college graduate and hence overstate the decline in research productivity. As an alternative, we redid all our results using two alternative deflators: first by adding 1 percent per year to the high-skilled nominal wage growth as a coarse adjustment and second using nominal GDP per person to deflate R&D expenditures, which according to the discussion surrounding equation (12) is a valid way to proceed. The results are shown in the online Appendix and are broadly similar, in part because the decreases in research productivity that we document are so large.

¹⁰The main surprise in confirming this observation is that w_t/A_{it} is constant. In particular, w_t is proportional to output per worker, and one might have expected that output per worker would grow with A_t (an average across varieties) but also with N_t , the number of varieties. However, this turns out not to be the case: Howitt includes a fixed factor of production (like land), and this fixed factor effectively eats up the gains from expanding variety. More precisely, the number of varieties grows with population while the amount of land per person declines with population, and these two effects exactly offset.

¹¹Building on the preceding footnote, it is worth also considering Peretto (2016b) in this context. Like Howitt, that paper has a fixed factor and for some parameter values, his setup also leads to constant research productivity. For other parameter values (e.g., if the fixed factor is turned off), the wage w_t grows both because of quality improvements and because of increases in variety. Nevertheless, deflating by the wage is still a good way to test the null hypothesis of endogenous growth: in that case, research productivity *rises* along an endogenous growth path. So the finding below that research productivity is declining is also relevant in this broader framework.

of research effort already incorporates a market-based adjustment for the depletion of talent: R&D spending weights workers according to their wage, and less talented researchers will naturally earn a lower salary. If more of these workers are hired over time, R&D spending will not rise by as much. Second, adjusting for λ only affects the magnitude of the trend in research productivity, not the overall qualitative fact of whether or not there is a downward trend. It is easy to deflate the growth rate of research effort by any particular value of λ to get a sense for how this matters; cutting our growth rates in half, an extreme adjustment, would still leave the nature of our results unchanged. Finally, there is no consensus on what value of λ one should use: Kremer (1993) even considers the possibility that it might be larger than one because of network effects and Zeira (2011) shows how patent races and duplication can occur even with $\lambda = 1$. Nevertheless, the Appendix shows the robustness of the main results in the paper to our baseline assumption by considering the case of $\lambda = 3/4$.

The remainder of the paper applies this framework in a wide range of different contexts: Moore's Law for semiconductors, agricultural crop yields, pharmaceutical innovation and mortality, and then finally at the firm level using Compustat data.

III. Moore's Law

One of the key drivers of economic growth during the last half century is Moore's Law: the empirical regularity that the number of transistors packed onto an integrated circuit serving as the central processing unit for a computer doubles approximately every two years. Figure 3 shows this regularity back to 1971. The log scale of this figure indicates the overall stability of the relationship, dating back nearly 50 years, as well as the tremendous rate of growth that is implied. Related formulations of Moore's Law involving computing performance per watt of electricity or the cost of information technology could also be considered, but the transistor count on an integrated circuit is the original and most famous version of the law, so we use that one here.

A doubling time of two years is equivalent to a constant exponential growth rate of 35 percent per year. We therefore measure the output of the idea production for Moore's Law as a stable 35 percent per year. Other alternatives are possible. For example, we could use decadal growth rates or other averages, and some of these approaches will be employed later in the paper. However, from the standpoint of understanding steady, rapid exponential growth for nearly half a century, the stability implied by the straight line in Figure 3 is a good place to start. And any slowing of Moore's Law would only reinforce the finding we are about to document.¹²

If the output side of Moore's Law is constant exponential growth, what is happening on the input side? Many commentators note that Moore's Law is not a law of nature but instead results from intense research effort: doubling the transistor density is often viewed as a goal or target for research programs. We measure research effort by deflating the nominal semiconductor R&D expenditures of all the

¹²For example, there is a recent shift away from speed and toward energy-saving features; see Flamm (2017) and Pillai (2016). However, our analysis still applies historically.

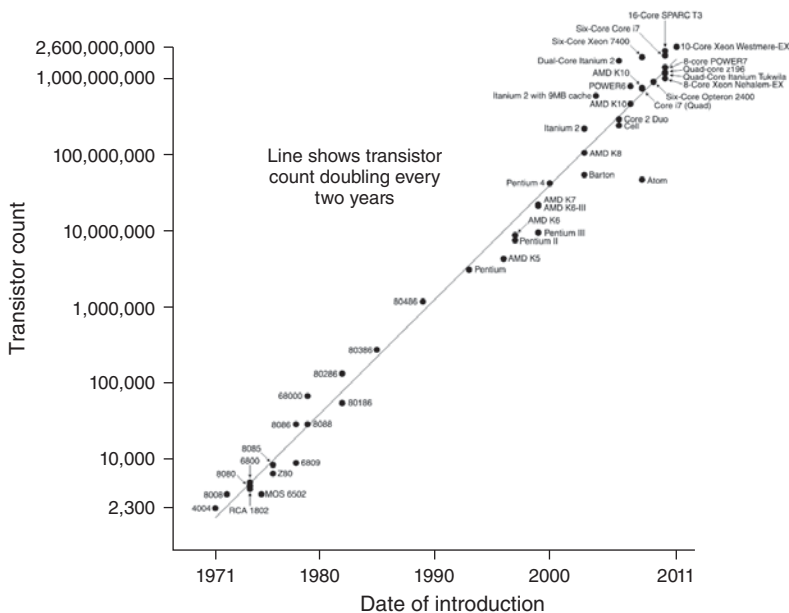


FIGURE 3. THE STEADY EXPONENTIAL GROWTH OF MOORE’S LAW

Source: Wikipedia (2017)

main firms by the nominal wage of high-skilled workers, as discussed above. Our semiconductor R&D series includes research spending by Intel, Fairchild, National Semiconductor, Motorola, Texas Instruments, Samsung, and more than two dozen other semiconductor firms and equipment manufacturers. More details are provided in the notes to Table 1 and in the online Appendix.

The striking fact, shown in Figure 4, is that research effort has risen by a factor of 18 since 1971. This increase occurs while the growth rate of chip density is more or less stable: the constant exponential growth implied by Moore’s Law has been achieved only by a massive increase in the amount of resources devoted to pushing the frontier forward.

Assuming a constant growth rate for Moore’s Law, the implication is that research productivity has fallen by this same factor of 18, an average rate of 6.8 percent per year. If the null hypothesis of constant research productivity were correct, the growth rate underlying Moore’s Law should have increased by a factor of 18 as well. Instead, it was remarkably stable. Put differently, because of declining research productivity, it is *around 18 times harder* today to generate the exponential growth behind Moore’s Law than it was in 1971.

The top panel of Table 1 reports the robustness of this result to various assumptions about which R&D expenditures should be counted. No matter how we measure R&D spending, we see a large increase in effective research and a corresponding large decline in research productivity. Even by the most conservative measure in the table, research productivity falls by a factor of 8 between 1971 and 2014.

The bottom panel of Table 1 considers an alternative to Moore’s Law as the “idea output” measure, focusing instead on TFP growth in the “semiconductor and related device manufacturing” industry (NAICS 334413) from the

TABLE 1—RESEARCH PRODUCTIVITY FOR MOORE'S LAW

	Factor decrease	Average growth (%)	Implied half-life (years)
<i>Moore's Law, 1971–2014</i>			
Baseline	18	−6.8	10.3
1. Narrow	8	−4.8	14.5
2. Narrow (downweight conglomerates)	11	−5.6	12.3
3. Broad (downweight conglomerates)	26	−7.6	9.1
4. Intel only (narrow)	347	−13.6	5.1
5. Intel + AMD (narrow)	352	−13.6	5.1
<i>TFP Growth in Semiconductors, 1975–2011</i>			
6. Narrow (no equipment R&D)	5	−3.2	21.4
7. Narrow (with equipment R&D)	7	−4.4	15.8
8. Broad (no equipment R&D)	11	−5.6	12.3
9. Broad (with equipment R&D)	13	−6.1	11.3

Notes: Research productivity is the ratio of idea output, either a constant 35 percent per year for the first panel or TFP growth in semiconductors for the second, to the effective number of researchers. The effective number of researchers is measured by deflating the nominal semiconductor R&D expenditures of key firms by the average wage of high-skilled workers. The R&D measures are based on Compustat (2016) data and PATSTAT data, assembled with assistance from Unni Pillai and Antoine Dechezlepretre. We start with the R&D spending data on 30 semiconductor firms plus an additional 11 semiconductor equipment manufacturers from all over the world. Next, we gathered data from PATSTAT on patents from the US patent office. The different rows in this table differ in how we add up the data across firms. There are two basic ways we treat the R&D data. In the *Narrow* treatment, we recognize that firms engage in different kinds of R&D, only some of which may be relevant for Moore's Law. We therefore weight a firm's R&D according to a (moving average) of the share of its patents that are in semiconductors (IPC group "H01L"). For example, in 1970, 75 percent of Intel's patents were for semiconductors, but by 2010 this number had fallen to just 8 percent. In the *Broad* category, we include all R&D by semiconductor firms like Intel and National Semiconductor but use the patent data to infer semiconductor R&D for conglomerates like IBM, RCA, Texas Instruments, Toshiba, and Samsung. The *downweight conglomerates* label means that we further downweight the R&D spending of conglomerates and newer firms like Micron and SK Hynix that focus on memory chips or chips for HDTVs and automobiles by a factor of 1/2, reflecting the possibility that even their semiconductor patenting data may be broader than Moore's Law. Rows 4 and 5 show results when we consider the *Narrow* measure of R&D but focus on only one or two firms. Rows 6 through 9 undertake the calculation using TFP growth in the "semiconductor and related device manufacturing" industry (334413) from the NBER/CES Manufacturing Industry Database; see Bartelsman and Gray (1996). We smooth TFP growth using the HP filter and lag R&D by 5 years in this calculation. In addition to the narrow/broad split, we also include and exclude R&D from semiconductor equipment manufacturers in this calculation (equipment is captured in a separate 6-digit industry, but there may be spillovers). See the online Appendix for more details. The implied half life is the number of years that it takes research productivity to fall in half at the measured growth rate.

NBER/CES Manufacturing Industry Database. Because of an acceleration in average TFP growth from around 8 percent per year to around 20 percent per year in the late 1990s, approximately a three-fold increase in the growth rate, these calculations show a slightly smaller decline in research productivity. Still, though, the increases in research effort are much larger than the increase in TFP growth, so research productivity falls substantially even using this alternative.

Caveats

Now is a good time to consider what could go wrong in our research productivity calculation at the micro level. Mismeasurement on both the output and input sides are clearly a cause for concern in general. However, there are two specific measurement problems that are worth considering in more detail. First, suppose there are "spillovers" from other sectors into the production of new ideas related to semiconductors. For example, progress in a completely different branch of materials




FIGURE 4. DATA ON MOORE'S LAW

Notes: The effective number of researchers is measured by deflating the nominal semiconductor R&D expenditures of key firms by the average wage of high-skilled workers and is normalized to 1 in 1970. The R&D data include research by Intel, Fairchild, National Semiconductor, Texas Instruments, Motorola, and more than two dozen other semiconductor firms and equipment manufacturers; see Table 1 for more details.

science may lead to a new idea that improves computer chips. Such positive spillovers are not a problem for our analysis; instead, they are one possible factor that our research productivity measure captures. Of course, other things equal, positive spillovers would show up as an *increase* in research productivity rather than as the declines that we document in this paper. Alternatively, if such spillovers were larger at the start of our time period than at the end, then this would be one possible story for why research productivity has declined.¹³

A type of measurement error that could cause our findings to be misleading is if we systematically understate R&D in early years and this bias gets corrected over time. In the case of Moore's Law, we are careful to include research spending by firms that are no longer household names, like Fairchild Camera and Instrument (later Fairchild Semiconductor) and National Semiconductor so as to minimize this bias: for example, in 1971, Intel's R&D was just 0.4 percent of our estimate for total semiconductor R&D in that year. Throughout the paper, we try to be as careful as we can with measurement issues, but this type of problem must be acknowledged.

IV. Agricultural Crop Yields

Our next application for measuring research productivity is agriculture. Due partly to the sector's historical importance, crop yields and agricultural R&D spending are relatively well measured. We begin in Figure 5 by showing research productivity for the agriculture sector as a whole. As our "idea output" measure, we use (a smoothed

¹³Lucking, Bloom, and Van Reenen (2017) provides an analysis of R&D spillovers using US firm-level data over the last three decades. They find evidence that knowledge spillovers are substantial, but have been broadly stable over time.




FIGURE 5. TFP GROWTH AND RESEARCH EFFORT IN AGRICULTURE

Notes: The effective number of researchers is measured by deflating nominal R&D expenditures by the average wage of high-skilled workers. Both TFP growth and US R&D spending (public and private) for the agriculture sector as a whole are taken from the US Department of Agriculture Economic Research Service (2018a,b). The TFP series is smoothed with an HP filter. Global R&D spending for agriculture is taken from Fuglie et al. (2011), Beintema et al. (2012), and Pardey et al. (2016).

version of) total factor productivity growth over the next five years. TFP growth declines slightly in agriculture, while effective research rises by about a factor of two between 1970 and 2007. Research productivity therefore declines over this period by a factor of nearly four, or at an average annual rate of 3.7 percent per year.

We now turn to the main focus of this section, research productivity for various agricultural crops. Ideally, we would use total factor productivity by crop as our idea output measure. Unfortunately, such a measure is not available because farms are “multiple input, multiple output” enterprises in which capital, labor, materials, and energy inputs are not easily allocated to individual crops. Instead, we use the growth rate of yield per acre as our measure of idea output. For the agriculture sector as a whole, growth in yield per acre and total factor productivity look similar.

For each of corn, soybeans, cotton, and wheat, we obtain data on both crop yields and R&D expenditures directed at improving crop yields. Figure 6 shows the annualized average 5-year growth rate of yields (after smoothing to remove shocks mostly due to weather). Yield growth has averaged around 1.5 percent per year since 1960 for these four crops, but with ample heterogeneity. These 5-year growth rates serve as our measure of idea output in studying the idea production function for seed yields.

The green lines in Figure 6 show measures of the “effective” number of researchers focused on each crop, measured as the sum of public and private R&D spending deflated by the wage of high-skilled workers. Two measures are presented. The faster-rising number corresponds to research targeted only at so-called biological efficiency. This includes cross-breeding (hybridization) and genetic modification directed at increasing yields, both directly and indirectly via improving insect resistance, herbicide tolerance, and efficiency of nutrient uptake, for example. The slower-growing number additionally includes research on crop protection

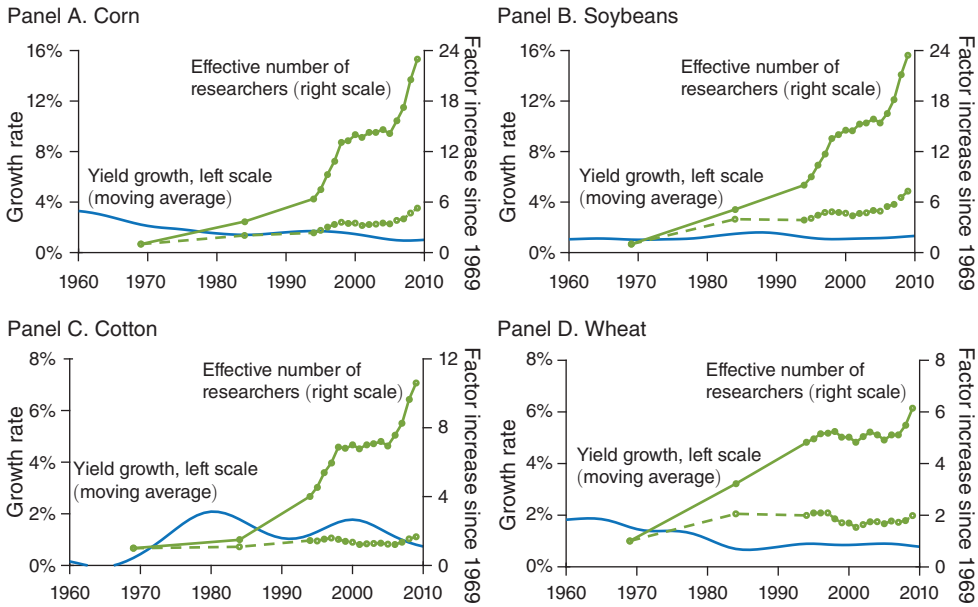


FIGURE 6. YIELD GROWTH AND RESEARCH EFFORT BY CROP

Notes: The blue line is the annual growth rate of the smoothed crop yields over the following 5 years; national realized yields for each crop are taken from the US Department of Agriculture National Agricultural Statistics Service (2016). The two green lines report *effective research*: the solid line is based on R&D targeting seed efficiency only; the dashed lower line additionally includes research on crop protection. Both are normalized to 1 in 1969. R&D expenditures are deflated by a measure of the nominal wage for high-skilled workers. See the online Appendix for more details.

and maintenance, which includes the development of herbicides and pesticides. The effective number of researchers has grown sharply since 1969, rising by a factor that ranges from 3 to more than 25, depending on the crop and the research measure.¹⁴

It is immediately evident from Figure 6 that research productivity has fallen sharply for agricultural yields: yield growth is relatively stable or even declining, while the effective research that has driven this yield growth has risen tremendously. Research productivity is simply the ratio of average yield growth divided by the number of researchers.

Table 2 summarizes the research productivity calculation for seed yields. As already noted, the effective number of researchers working to improve seed yields rose enormously between 1969 and 2009. For example, the increase was more than a factor of 23 for both corn and soybeans when research is limited to seed efficiency. If yield growth were constant (which is not a bad approximation across the four crops as shown in Figure 6), then research productivity would on average

¹⁴Our measure of R&D inputs consists of the sum of R&D spending by the public and private sectors in the United States. Data on private sector biological efficiency and crop protection R&D expenditures are from an updated USDA series based on Fuglie et al. (2011), with the distribution of expenditure by crop taken from Perrin, Kunnings, and Ihnen (1983); Fernandez-Cornejo et al. (2004); Traxler et al. (2005); Huffman and Evenson (2006); and Centre for Industry Education Collaboration (2016). Data on US public sector R&D expenditure by crop are from the US Department of Agriculture National Institute of Food and Agriculture Current Research Information System (2016) and Huffman and Evenson (2006), with the distribution of expenditure by research focus taken from Huffman and Evenson (2006).

TABLE 2—RESEARCH PRODUCTIVITY IN AGRICULTURE, 1969–2009

	Effective research		Research productivity	
	Factor increase	Avg. growth (%)	Factor decrease	Avg. growth (%)
Research on seed efficiency only				
Corn	23.0	7.8	52.2	−9.9
Soybeans	23.4	7.9	18.7	−7.3
Cotton	10.6	5.9	3.8	−3.4
Wheat	6.1	4.5	11.7	−6.1
Research includes crop protection				
Corn	5.3	4.2	12.0	−6.2
Soybeans	7.3	5.0	5.8	−4.4
Cotton	1.7	1.3	0.6	+1.3
Wheat	2.0	1.7	3.8	−3.3
Agriculture				
US research, 1970–2008	1.9	1.8	3.9	−3.7
Global research, 1980–2010	1.6	1.6	5.2	−5.5

Notes: Research productivity is the ratio of idea output, yield growth, to the effective number of researchers, measured as R&D expenditures deflated by the nominal wage for high-skilled workers. In the first panel of results, the research input is based on R&D expenditures for seed efficiency only. The second panel additionally includes research on crop protection. See the online Appendix for more details.

decline by this same factor. The last 2 columns of Table 2 show this to be the case. On average, research productivity declines for crop yields by about 6 percent per year using the narrow definition of research and by about 4 percent per year using the broader definition.

A potential source of mismeasurement for this case relates to the quality of land inputs. What if researchers are devoting their efforts to bringing lower-quality land into production? This could show up as a decrease in average yields, even as research is increasing yields for any given quality of land. First, at a high level, it's worth noting that the total acreage of cotton and wheat planted in the United States has declined over our time period. (By contrast, acreage devoted to soybeans has doubled, while that for corn has increased slightly.) The declining acreage for cotton and wheat, and roughly constant acreage for corn, does not suggest on the face of it that changes on the extensive margin for those crops are crucial. Moreover, in order for what we find to be consistent with constant research productivity, we'd need average yields in the absence of research to be falling (due to lower land quality, say) at a rate that is growing exponentially in magnitude. This also seems unlikely. As an additional robustness check, we calculated our measures of research productivity using state-level estimates of seed yield growth as our idea output measure (maintaining our broad research effort measure, since the nonrivalry of ideas means that research everywhere could be relevant for seed yields in each state). We found that, for each crop, the vast majority of individual states experienced declines in research productivity over our time period, mirroring the national-level results.

V. Mortality and Life Expectancy

Health expenditures account for around 18 percent of US GDP, and a healthy life is one of the most important goods we purchase. Our third collection of case studies examines the productivity of medical research.

A. *New Molecular Entities*

Our first example from the medical sector is a fact that is well known in the literature, recast in terms of our research productivity calculation. New molecular entities (NMEs) are novel compounds that form the basis of new drugs. Historically, the number of NMEs approved by the Food and Drug Administration each year shows little or no trend, while the number of dollars spent on pharmaceutical research has grown dramatically; for example, see Akcigit and Liu (2016). We reexamine this fact using our measure of research productivity, i.e., deflating pharmaceutical research by the high-skilled wage. The details of this calculation are reported in the online Appendix. The result is that research effort rises by a factor of 9, while research productivity falls by a factor of 11 by 2007 before rising in recent years so that the overall decline by 2014 is a factor of 5. Over the entire period, research effort rises at an annual rate of 6.0 percent, while research productivity falls at an annual rate of 3.5 percent. Of course, it is far from obvious that simple counts of NMEs appropriately measure the output of ideas; we would really like to know how important each innovation is. In addition, NMEs suffer from an important aggregation issue, adding up across a wide range of health conditions. These limitations motivate the approach described next, where we turn to the productivity of medical research for specific diseases.

B. *Years of Life Saved*

To measure idea output in treating diseases, we begin with life expectancy. Figure 7 shows US life expectancy at birth and at age 65. This graph makes the important point that life expectancy is one of the few economic goods that does not exhibit exponential growth. Instead, arithmetic growth is a better description. Since 1950, US life expectancy at birth has increased at a relatively stable rate of 1.8 years each decade, and life expectancy at age 65 has risen at 0.9 years per decade. Linear increases in life expectancy seem to coincide with stable economic growth.¹⁵

Also shown in the graph is the well-known fact that overall life expectancy grew even more rapidly during the first half of the twentieth century, at around 3.8 years per decade. This raises the question of whether even arithmetic growth is an appropriate characterization. We believe that it is for two reasons. First, there is no sign of a slowdown in the years gained per decade since 1950, either in life expectancy at birth or in life expectancy at age 65. The second reason is a fascinating empirical regularity documented by Oeppen and Vaupel (2002). That paper shows that “record female life expectancy,” the life expectancy of women in the country for which they live the longest, has risen at a remarkably steady rate of 2.4 years per decade ever since 1840. Steady linear increases in life expectancy, not exponential ones, seem to be the norm.

For this reason, we use “years of life saved,” that is, the change in life expectancy rather than its growth rate, as a measure of idea output. Because the growth rate of life expectancy is declining, our results would be even stronger under that alternative.

¹⁵For example, see Nordhaus (2003), Hall and Jones (2007), and Dalgaard and Strulik (2014).




FIGURE 7. UNITED STATES LIFE EXPECTANCY

Sources: National Center for Health Statistics (2014) and Clio Infra (2016)

C. Years of Life Saved from Specific Diseases

To measure the years of life saved from reductions in disease-specific mortality, consider a person who faces two age-invariant Poisson processes for dying, with arrival rates δ_1 and δ_2 . We think of δ_1 as reflecting a particular disease we are studying, such as cancer or heart disease, and δ_2 as capturing all other sources of mortality. The probability a person lives for at least x years before succumbing to type i mortality is the survival rate $S_i(x) = e^{-\delta_i x}$, and the probability the person lives for at least x years before dying from any cause is $S(x) = S_1(x) S_2(x) = e^{-(\delta_1 + \delta_2)x}$. Life expectancy at age a , $LE(a)$ is then well known to equal

$$(15) \quad LE(a) = \int_0^\infty S(x) dx = \int_0^\infty e^{-(\delta_1 + \delta_2)x} dx = \frac{1}{\delta_1 + \delta_2}.$$

Now consider how life expectancy changes if the type i mortality rate changes slightly. It is easy to show that the expected years of life saved by the mortality change is

$$(16) \quad dLE(a) = \frac{\delta_i}{\delta_1 + \delta_2} \cdot LE(a) \cdot \left(-\frac{d\delta_i}{\delta_i} \right).$$

That is, the expected years of life saved from a decline in, say, cancer mortality is the product of three terms. First is the fraction of deaths that result from cancer. Second is the average years of life lost if someone dies from cancer at age a , and the final term is the percentage decline in cancer mortality.¹⁶

¹⁶Our measures of life expectancy and mortality from all sources by age come from the Human Mortality Database (2016) (<http://mortality.org>). To measure the percentage declines in mortality rates from cancer, we use the age-adjusted mortality rates for people ages 50 and over computed from 5-year survival rates, taken from the National Cancer Institute’s Surveillance, Epidemiology, and End Results program (<http://seer.cancer.gov/>).

Vaupel and Canudas Romo (2003) shows that this expression generalizes to a much richer setting. In particular, the expected years of life saved is the product of three terms with the same interpretation. For example, they allow for an arbitrary number of causes of death each of which has a mortality rate that varies arbitrarily with age.¹⁷

The research input aimed at reducing mortality from a given disease is at first blush harder to measure. For example, it is difficult to get research spending broken down into spending on various diseases. Nevertheless, we implement a potential solution to this problem by measuring the number of scientific publications in PubMed that have “Neoplasms,” for example, as a MeSH (Medical Subject Heading) term. MeSH is the National Library of Medicine’s controlled vocabulary thesaurus.¹⁸ We do this in two ways. Our broader approach (publications) uses all publications with the appropriate MeSH keyword as our input measure. Our narrower approach (trials) further restricts our measure to those publications that according to MeSH correspond to a clinical trial. Rather than using scientific publications as an *output* measure, as other studies have done, we use publications and clinical trials as *input* measures to capture research effort aimed at reducing mortality for a particular disease.¹⁹

Figure 8 shows our basic “idea output” and “idea input” measures for mortality from all cancers, from breast cancer, and from heart disease. Heart disease and cancer are the top two leading causes of death in the United States, and in the spirit of looking as narrowly as possible, we also chose to look at breast cancer mortality. For the two cancer types, we use the 5-year mortality rate conditional on being diagnosed with either type of cancer and see an S-shaped decline since 1975. This translates into a hump-shaped *Years of life saved per 1,000 people*, the empirical analog of equation (16). For example, for all cancers, the years of life saved series peaks around 1990 at more than 100 years of life saved per 1,000 people before declining to around 60 years in the 2000s. For heart disease, a substantial part of the decline in deaths comes from people not contracting the disease in the first place, so we focus on the (smoothed) crude death rate for people aged 55 to 64. The death rate declines at different rates in different periods, leading to a series of humps in years of life saved, but overall there is no large trend in this measure of idea output.

The right panels of Figure 8 show our research input measure based on PubMed publication statistics. Total publications for all cancers increased by a factor of 3.5 between 1975 and 2006 (the years for which we’ll be able to compute research productivity), while publications restricted to clinical trials increased by a factor of

¹⁷Their formula involves an extra covariance term as well. In particular, the covariance between the age-specific percentage decline in mortality associated with cancer and the years of life saved at age a when cancer is averted. When the percentage decline in mortality rates is the same across ages, this covariance is zero. More generally, it can differ from zero, but in many of the calculations in their paper, the covariance is small.

¹⁸For more information on MeSH, see <https://www.nlm.nih.gov/mesh/>. Our queries of the PubMed data use the webtool created by the Institute for Biostatistics and Medical Informatics (IBMI) Medical Faculty, University of Ljubljana, Slovenia (available at <http://webtools.mf.uni-lj.si/>).

¹⁹In independent work, Lichtenberg (2017) takes a similar approach in an econometric framework for the years 1999–2013. He uses a difference-in-differences specification to document an economically-significant correlation between research publications related to various cancer sites and subsequent mortality and years of life saved. Lichtenberg (2018) extends this approach further back in time to the period 1946–2015 and continues to find a relationship between publications and 5-year survival rates.

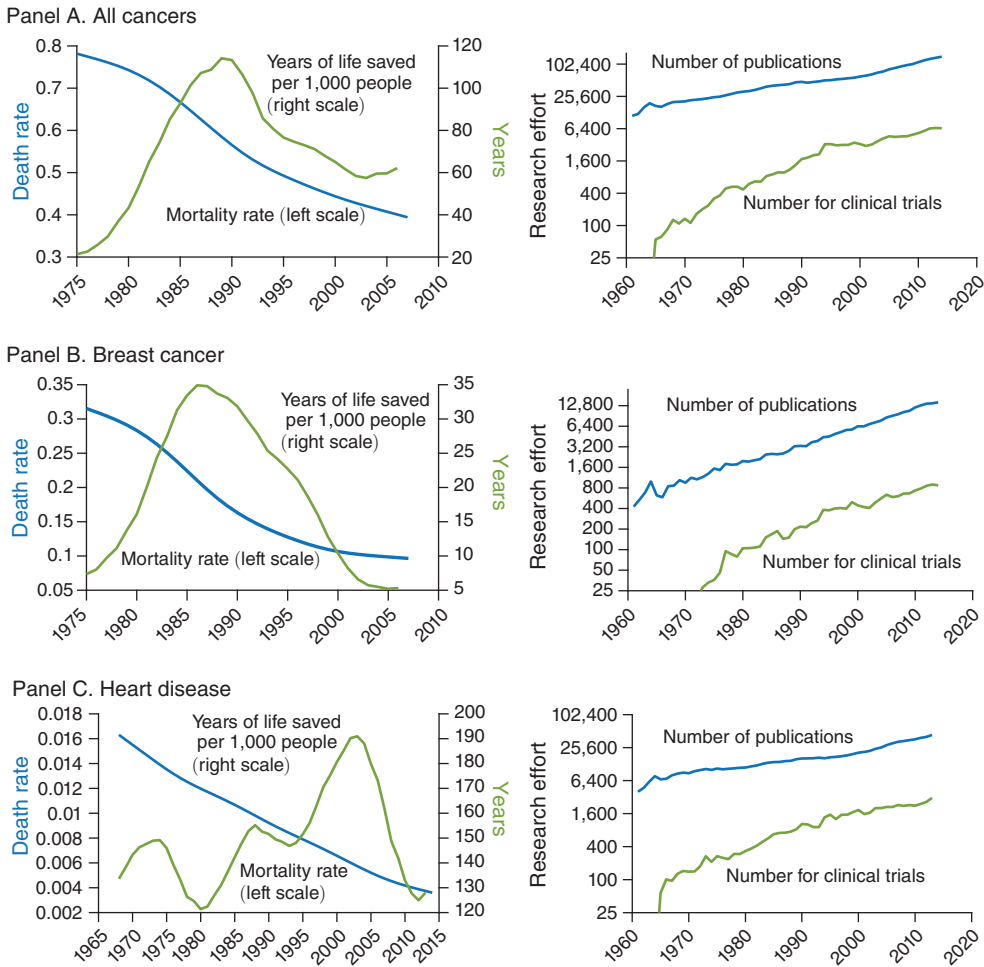


FIGURE 8. MORTALITY, YEARS OF LIFE SAVED, AND RESEARCH EFFORT

Notes: For the two cancer panels, the mortality rate is computed as negative the log of the (smoothed) five-year survival rate for cancer for people ages 50 and higher, from the National Cancer Institute’s Surveillance, Epidemiology, and End Results program at <http://seer.cancer.gov/>. For heart disease, we report the crude death rate in each year for people aged 55–64. *Years of life saved per 1,000 people* is computed using equation (16), as described in the text. Research effort is measured by the number of related publications and clinical trials, taken from the PubMed publications database. For *publications*, the research input is based on all publications in PubMed with “Neoplasms” or “Breast Neoplasms” or “Heart Diseases” as a MeSH keyword. The lines for “clinical trials” restrict further to publications involving clinical trials.

14.1 during this same period. A similar pattern is seen for research on breast cancer and heart disease.

Research productivity for our medical research applications is computed as the ratio of years of life saved to the number of publications. Figure 9 shows our research productivity measures. The hump shape present in the years-of-life-saved measure carries over here. Research productivity rises until the mid-1980s and then falls. Overall, between 1975 and 2006, research productivity for all cancers declines by a factor of 1.2 using all publications and a factor of 4.8 using clinical trials. The declines for breast cancer and heart disease are even larger, as shown in Table 3.




FIGURE 9. RESEARCH PRODUCTIVITY FOR MEDICAL RESEARCH

Note: Research productivity is computed as the ratio of years of life saved to the number of publications.

TABLE 3—RESEARCH PRODUCTIVITY FOR MEDICAL RESEARCH

Disease	Effective research		Research productivity	
	Factor increase	Avg. growth (%)	Factor decrease	Avg. growth (%)
All publications				
Cancer, all types	3.5	4.0	1.2	-0.6
Breast cancer	5.9	5.7	8.2	-6.8
Heart disease	5.1	3.6	5.3	-3.7
Clinical trials only				
Cancer, all types	14.1	8.5	4.8	-5.1
Breast cancer	16.3	9.0	22.6	-10.1
Heart disease	24.2	7.1	25.3	-7.2

Notes: Research productivity is computed as the ratio of years of life saved to the number of publications. In the first panel, the research input is based on all publications in PubMed with “Neoplasms” or “Breast Neoplasms” or “Heart Diseases” as a MeSH keyword. The second panel restricts to only publications involving clinical trials. Results for cancer and breast cancer cover the years 1975–2006, while those for heart disease apply to 1968–2011. See the online Appendix for more details.

Several general comments about research productivity for medical research deserve mention. First, for this application, the units of research productivity are different than what we’ve seen so far. For example, between 1985 and 2006, declining research productivity means that the number of years of life saved per 100,000 people in the population by each publication of a clinical trial related to cancer declined from more than 8 years to just over one year. For breast cancer, the changes are even starker: from around 16 years per clinical trial in the mid-1980s to less than 1 year by 2006.

Next, however, notice that the changes were not monotonic if we go back to 1975. Between 1975 and the mid-1980s, research productivity for these two cancer research categories increased quite substantially. The production function for new ideas is obviously complicated and heterogeneous. These cases suggest that it may get easier to find new ideas at first before getting harder, at least in some areas.

VI. Research Productivity in Firm-Level Data

Our studies of semiconductors, crops, and health are illuminating, but at the end of the day, they are just case studies. One naturally wonders how representative they are of the broader economy. In addition, some growth models associate each firm with a different variety: perhaps the number of firms making corn or semiconductor chips is rising sharply, so that research effort per firm is actually constant, as is research productivity at the firm level. Declining research productivity for corn or semiconductors could in this view simply reflect a further composition bias.²⁰

To help address these concerns, we report two sets of results with firm-level data. Our first set is Compustat (2016) data on US publicly-traded firms. Our second set is administrative data from the Census of Manufacturing. Each dataset has its strengths. Compustat includes a longer time series as well as firms from outside manufacturing. The Census covers the universe of manufacturing firms rather than just those that are publicly traded.

The strength of the firm-level data is that they are more representative than the case studies, but of course they too have limitations. Publicly-traded firms and manufacturing firms are each still a selected sample, and our measures of “ideas” and research inputs are likely to be less precise. And in some models, the product line rather than the firm is the right unit of observation. Creative destruction may make the firm-level results harder to interpret. This is because part of the effect of a firm’s own R&D on its growth rate may reflect not just the impact on its productivity, but also a gain in market share at the expense of another firm (business stealing). Fortunately, empirical evaluations of the magnitude of R&D-induced business stealing using Compustat data find it to be quantitatively dominated by the knowledge creating effects of R&D (see Bloom, Schankerman, and Van Reenen 2013 and Lucking, Bloom, and Van Reenen 2017). With these caveats in mind, we view this firm-level evidence as a helpful complement to the case studies.

A. Compustat Results

As a measure of the output of the idea production function, we use decadal averages of annual growth in sales revenue, market capitalization, employment, and revenue labor productivity within each firm. We take the decade as our period of observation to smooth out fluctuations.

Why would growth in sales revenue, market cap, or employment be informative about a firm’s production of ideas? This approach follows a recent literature

²⁰For example, Peretto (1998, 2016b) emphasize this perspective on varieties, while Aghion and Howitt (1992) takes the alternative view that different firms may be involved in producing the same variety. Klette and Kortum (2004) allows the number of varieties produced by each firm to be heterogeneous and to evolve over time.

emphasizing precisely these links. Many papers have shown that news of patent grants for a firm has a large immediate effect on the firm's stock market capitalization (e.g., Blundell, Griffith, and Van Reenen 1999; Kogan et al. 2015). Patents are also positively correlated with the firm's subsequent growth in employment and sales.

More generally, in models in the tradition of Lucas (1978), Hopenhayn (1992), and Melitz (2003), increases in the fundamental productivity of a firm show up in the long run as increases in sales and firm size, but *not* as increases in sales revenue per worker.²¹ This motivates our use of sales revenue or employment to measure fundamental productivity. Hsieh and Klenow (2009) and Garcia-Macia, Hsieh, and Klenow (2016) are recent examples of papers that follow a related approach. Of course, in more general models with fixed overhead labor costs for example, revenue labor productivity (i.e., sales revenue per worker) and TFPR can be related to fundamental productivity (e.g., Bartelsman, Haltiwanger, and Scarpetta 2013). And sales revenue and employment can change for reasons other than the discovery of new ideas. We try to address these issues by also looking at revenue productivity and through various sample selection procedures, discussed below. These problems also motivate the earlier approach of looking at case studies.

To measure the research input, we use a firm's spending on research and development from Compustat. This means we are restricted to publicly-listed firms that report formal R&D, and such firms are well-known to be a select sample (e.g., disproportionately in manufacturing and large). We look at firms since 1980 that report nonzero R&D, and this restricts us to an initial sample of 15,128 firms. Our additional requirements for sample selection in our baseline sample are:

- (i) We observe at least 3 annual growth observations for the firm in a given decade. These growth rates are averaged to form the idea output growth measure for that firm in that decade.
- (ii) We only consider decades in which our idea output growth measure for the firm is positive (negative growth is clearly not the result of the firm innovating).
- (iii) We require the firm to be observed (for both the output growth measure and the research input measure) for two consecutive decades. Our decades are the 1980s, the 1990s, the 2000s (which refers to the 2000–2007 period), and the 2010s (which refers to the 2010–2015 period); we drop the years 2008 and 2009 because of the financial crisis.

We relax many of these conditions in our robustness checks.

²¹This is obvious when one thinks about the equilibrium condition for the allocation of labor across firms in simple settings: in equilibrium, a worker must be indifferent between working in two different firms, which equalizes wages. But wages are typically proportional to output per worker. Moreover, with Cobb-Douglas production and a common exponent on labor, sales revenue per worker would be precisely equated across firms even if they had different underlying productivities. In a Lucas (1978) span of control setting, more productive firms just hire more workers, which drives down the marginal product until it is equated across firms. In alternative settings with monopolistic competition, it is the price of a particular variety that declines as the firm expands. Regardless, higher fundamental productivity shows up as higher employment or sales revenue, but not in higher sales per employee.

TABLE 4—RESEARCH PRODUCTIVITY IN COMPUSTAT FIRM-LEVEL DATA

Sample	Effective research		Research productivity	
	Factor increase	Avg. growth (%)	Factor decrease	Avg. growth (%)
Sales revenue				
2 decades (1,712 firms)	2.0	6.8	3.9	-13.6
3 decades (469 firms)	3.8	6.7	9.2	-11.1
4 decades (149 firms)	13.7	8.7	40.3	-12.3
Market cap				
2 decades (1,124 firms)	2.2	8.0	3.4	-12.2
3 decades (335 firms)	3.1	5.6	6.3	-9.2
4 decades (125 firms)	7.9	6.9	14.0	-8.8
Employment				
2 decades (1,395 firms)	2.2	8.0	2.8	-10.3
3 decades (319 firms)	4.0	6.9	18.2	-14.5
4 decades (101 firms)	13.9	8.8	31.5	-11.5
Revenue labor productivity				
2 decades (1,444 firms)	1.9	6.4	2.2	-7.9
3 decades (337 firms)	1.6	2.4	2.5	-4.5
4 decades (109 firms)	2.5	3.1	3.5	-4.2

Notes: The table shows averages of firm-level outcomes for effective research and research productivity. *Sales revenue* and *Market cap* are deflated by the GDP implicit price deflator. *Revenue labor productivity* is deflated *Sales revenue* divided by *Employment*. The effective number of researchers (*Effective research*) is measured by deflating nominal R&D expenditures by the average wage of high-skilled workers. The average growth rate across two decades is computed by dividing by ten years (e.g., between 1985 and 1995); others follow this same approach. Averages are computed by weighting firms by the median number of effective researchers in each firm across the decades. See the online Appendix for more details.

Table 4 shows our research productivity calculation for various cuts of the Compustat data. In all samples, there is substantial growth in the effective number of researchers within each firm, with growth rates averaging between 2.4 percent and 8.8 percent per year. Under our null hypothesis, this rapid growth in research should translate into higher growth rates of firm-level sales and employment with a constant level of research productivity. Instead, what we see in Table 4 are rapid declines in firm-level research productivity across all samples, at growth rates that range from -4.2 percent to -14.5 percent per year for multiple decades.

Averaging across all our samples, research productivity falls at a rate of about 9 percent per year, cumulating to a 2.5-fold decline every decade. At this rate, research productivity declines by a factor of about 15 over three decades of changes; put differently, it requires 15 times more researchers today than it did 30 years ago to produce the same rate of firm revenue growth.

Figure 10 demonstrates the heterogeneity across firms in our Compustat sample by showing the distribution of the factor changes in effective research and research productivity across firms. In this figure, we focus on the results for sales revenue for firms observed for two decades, but the results with other output measures and other time horizons are similar; additional results are available in the online Appendix.

The heterogeneity across firms is impressive and somewhat reminiscent of the heterogeneity we see in our case studies. Nevertheless, it is clear from the histogram that there is essentially no evidence that constant research productivity is

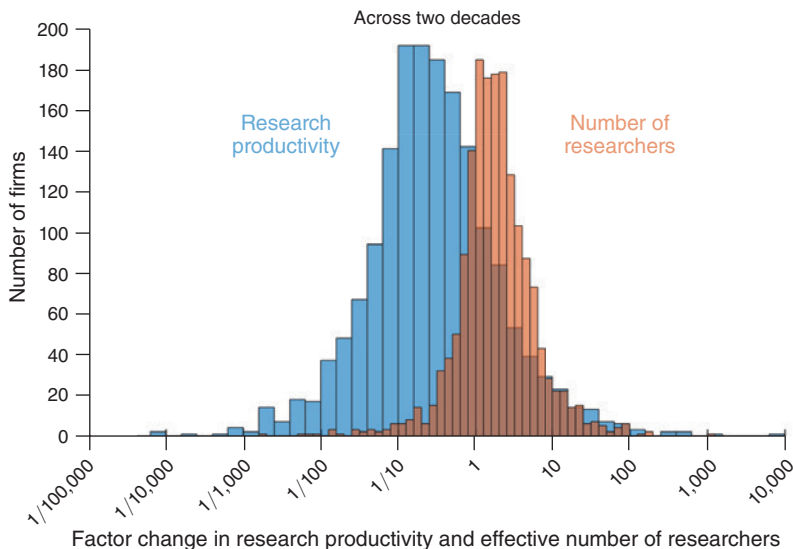


FIGURE 10. COMPUSTAT DISTRIBUTIONS, SALES REVENUE (TWO DECADES)

Notes: Based on 1,712 firms. 22.1 percent of firms have increasing research productivity. Only 3.0 percent have research productivity that is roughly constant, defined as a growth rate whose absolute value is less than 1 percent per year. See Notes to Table 4 for more details.

a good characterization of the firm-level data. The average, median, and modal firms experience large declines in research productivity. There is a long tail of firms experiencing even larger declines but also a small minority of firms that see increases in research productivity. The fraction of firms that exhibit something like constant research productivity is tiny. For example, less than 5 percent of firms across our three time frames have research productivity changing (either rising or falling) by less than 1 percent per year on average.

Table 5 provides additional evidence of robustness. In the interests of brevity, we report these results for the sales revenue and for firms that we observe across three decades, but the results for our other output measures and time frames are similar. The first row repeats the benchmark results described earlier. The second and third rows relax the requirement that sales growth is positive. This increases the sample size considerably and brings the R&D growth numbers down substantially. In both cases, research productivity falls sharply, reassuring us that this sample selection criteria is not driving the results. The fourth row imposes the restriction that research is increasing across the observed decades. The fifth row tightens our restrictions and drops firm in which sales revenue declines on average in any decade. The sixth row uses median sales growth rather than mean as our output measure. The seventh row reports unweighted averages rather than weighting firms by the effective number of researchers. And the eighth row uses so-called “DHS growth rates” defined following Davis, Haltiwanger, and Schuh (1996); this bounds growth rates between -2 and $+2$, addressing any concerns over outliers. The general finding of substantial declines in research productivity is robust.

TABLE 5—COMPUSTAT SALES DATA ACROSS THREE DECADES: ROBUSTNESS

Case	Effective research		Research productivity	
	Factor increase	Avg. growth (%)	Factor decrease	Avg. growth (%)
Benchmark (469 firms)	3.8	6.7	9.2	−11.1
Winsorize $g < 0.01$ (986 firms)	2.3	4.1	7.9	−10.3
Winsorize top/bottom (986 firms)	2.3	4.1	6.0	−8.9
Research must increase (356 firms)	5.1	8.1	11.6	−12.3
Drop if any negative growth (367 firms)	5.6	8.6	17.9	−14.4
Median sales growth (586 firms)	3.8	6.6	6.3	−9.2
Unweighted averages (469 firms)	3.8	6.7	9.2	−11.1
DHS growth rates (470 firms)	3.8	6.7	3.6	−6.4

Notes: Robustness results reported for the sample of changes across three decades. *Winsorize* $g < 0.01$ means we replace any idea output measure that is less than 1 percent annually with a value of 1 percent. *Winsorize top/bottom* does this same thing but winsorizes an equal number of firms at the top of the idea output distribution. *Research must increase* means we require that the research measure be rising across the decades. *Drop if any negative growth* means we drop firms that have any decade (across our 1980–2015 period) in which average market cap growth is negative. *Median sales growth* uses the median of sales revenue growth in each decade rather than the mean. *Unweighted averages* gives each firm equal weight in computing summary statistics, rather than weighting each firm by its effective number of researchers. *DHS growth rates* uses $(y_t - y_{t-1}) / (0.5 \times y_t + 0.5 \times y_{t-1})$ to compute the growth rate of research productivity at the firm level to reduce problems with small denominators.

B. Census of Manufacturing Results

We also look at firms in the US Census of Manufacturing (Bureau of the Census 2015b), using data from the responses to mandatory survey responses on sales (Economic Census) and on R&D activity (Bureau of the Census 2009, 2015a, SIRD/BRDIS). These data have a number of differences with Compustat which makes them a valuable robustness check. First, they use the responses to official (mandatory) government surveys, rather than audited financial accounts. So this could potentially induce some measurement error (if the auditing process helps to eliminate recording errors for example), but it also helps to address concerns over potential bias in reported accounting data for publicly-listed firms (e.g., if firms manipulate their reported R&D activity to influence their stock market valuation). Second, they cover the activity of all firms operating in the United States, public and private, including the subsidiaries of foreign multinationals. Thus, smaller firms, start-ups, and subsidiaries of overseas firms are included.²² Third, the Census data exclude the R&D and sales activities of US firms abroad, which for large manufacturing firms is often substantial. By contrast, Compustat reports the global consolidated accounts, so overseas sales and R&D will be included in the totals.²³ Fourth, the Census also collects data on the number of scientists and engineers engaged in R&D activity, providing a quantity measure of innovation inputs. Finally, the Census compares the figures for large firms against administrative data, e.g., IRS tax returns and social security filings, helping to ensure data accuracy.

Our sample includes all firms that reported manufacturing shipments (sales) in the Economic Censuses (CMF) of 1982, 1992, 2002, and 2012, as well as positive

²²Including the R&D and performance of small firms is potentially important as many innovations may come from younger firms. For example, see Acs and Audretsch (1988).

²³The Compustat Segment data do provide some breakdowns by geography and sector, but these are not consistently reported over a long enough time period nor for enough firms (especially for R&D) for our exercise.

TABLE 6—CENSUS OF MANUFACTURING RESULTS, ACROSS TWO DECADES (1992–2002, 2002–2012)

Case	Effective research		Research productivity	
	Factor increase	Avg. growth (%)	Factor decrease	Avg. growth (%)
1. Benchmark	1.2	1.6	2.2	-7.8
2. Winsorize $g < 0.01$	1.2	1.6	1.9	-6.0
3. Winsorize top/bottom	1.2	1.6	1.7	-4.9
4. Unweighted	1.0	0.0	1.9	-8.1
5. Research = scientists	1.3	2.3	2.3	-6.0

Notes: Research productivity is the ten-year DHS growth in real sales divided by mean R&D spending, deflated by the skilled wage, over those ten years. Research productivity growth is then calculated as the percent change in research productivity compared to ten years earlier. In row 2, idea output (sales growth) is winsorized from below at 1 percent. In row 3, idea output (sales growth) is winsorized from below at 1 percent and from above such that an equal number of firms are winsorized in each tail. In row 4, the mean is unweighted. In row 5, the denominator in research productivity is the number of scientists and engineers. In rows 1 to 3, the mean of the growth rate of R&D is weighted by mean R&D over the past 20 years. In row 5, the mean of the growth rate of scientists and engineers is weighted by mean R&D over the past 20 years. *Factor decrease* is calculated as $1/(1 - \text{mean})$ where mean is the mean of the research productivity growth weighted by the average R&D spending over the past 20 years. *Average growth* is calculated as $1 - (1 - \text{mean})^{1/10}$ where mean is the mean of research productivity growth weighted by the average R&D spending over the past 20 years. The sample includes 1,300 firms and 2,700 observations for all cells.

R&D expenditure in the BRDIS (or SIRD before 2008) surveys of R&D in at least one year in each decade. The data span 1,300 firms over 2,700 observations (where numbers have been rounded for disclosure purposes).

Research productivity results in the Census data are reported in Table 6 and are similar to what we saw in the Compustat sample. In row 1 we see that while effective research inputs have risen by 1.6 percent a year on average research productivity has fallen by 7.8 percent between 1992 and 2012. Part of the explanation is that our early period includes the “new economy” while our later period includes the Great Recession; this is a limitation of the Census sample. In row 2 we winsorize the growth rate in sales from below at 1 percent, finding that research productivity now falls by 6 percent. In row 3 we winsorize both the top and the bottom of the distribution of sales growth by the same share and still find a substantial decline in research productivity of 4.9 percent. In row 4 rather than weight firms by mean R&D spending over the 20 year window we provide unweighted results showing a larger decline of -8.1 percent. Finally, in row 5 we use the number of research scientists and engineers that is collected in the BRDIS/SIRD survey form as our measure of scientific input. This addresses any concerns over the deflator for R&D inputs as we are using a quantity (number of scientists and engineers) input measure, and we find once again a significant decline in research productivity of 6 percent.

VII. Discussion

A. Summary and Semi-Endogenous Growth Theory

The evidence presented in this paper concerns the extent to which a constant level of research effort can generate constant exponential growth, either in the economy as a whole or within relatively narrow categories, such as a firm or a seed type or a health condition. We provide consistent evidence that the historical answer to this question is “no”: as summarized in Table 7, research productivity is declining at a substantial rate in virtually every place we look. The table also provides a way to

TABLE 7—SUMMARY OF THE EVIDENCE ON RESEARCH PRODUCTIVITY

Scope	Time period	Average annual growth rate (%)	Half-life (years)	Dynamic diminishing returns, β
Aggregate economy	1930–2015	−5.1	14	3.1
Moore's Law	1971–2014	−6.8	10	0.2
Semiconductor TFP growth	1975–2011	−5.6	12	0.4
Agriculture, US R&D	1970–2007	−3.7	19	2.2
Agriculture, global R&D	1980–2010	−5.5	13	3.3
Corn, version 1	1969–2009	−9.9	7	7.2
Corn, version 2	1969–2009	−6.2	11	4.5
Soybeans, version 1	1969–2009	−7.3	9	6.3
Soybeans, version 2	1969–2009	−4.4	16	3.8
Cotton, version 1	1969–2009	−3.4	21	2.5
Cotton, version 2	1969–2009	+1.3	−55	−0.9
Wheat, version 1	1969–2009	−6.1	11	6.8
Wheat, version 2	1969–2009	−3.3	21	3.7
New molecular entities	1970–2015	−3.5	20	...
Cancer (all), publications	1975–2006	−0.6	116	...
Cancer (all), trials	1975–2006	−5.7	12	...
Breast cancer, publications	1975–2006	−6.1	11	...
Breast cancer, trials	1975–2006	−10.1	7	...
Heart disease, publications	1968–2011	−3.7	19	...
Heart disease, trials	1968–2011	−7.2	10	...
Compustat, sales	3 decades	−11.1	6	1.1
Compustat, market cap	3 decades	−9.2	8	0.9
Compustat, employment	3 decades	−14.5	5	1.8
Compustat, sales/employment	3 decades	−4.5	15	1.1
Census of Manufacturing	1992–2012	−7.8	9	...

Notes: The growth rates of research productivity are taken from other tables in this paper. The half-life is the number of years it takes for research productivity to fall in half at this growth rate. The last column reports the extent of dynamic diminishing returns in producing exponential growth, according to equation (17). This measure is only reported for cases in which the idea output measure is an exponential growth rate (i.e., not for the health technologies, where units would matter).

quantify the magnitude of the declines in research productivity by reporting the half-life in each case. Taking the aggregate economy number as a representative example, research productivity declines at an average rate of 5.3 percent per year, meaning that it takes around 13 years for research productivity to fall by half. Or put another way, the economy has to double its research efforts every 13 years just to maintain the same overall rate of economic growth.

A natural question is whether these empirical patterns can be reproduced in a general equilibrium model of growth. One class of models that is broadly consistent with this evidence is the semi-endogenous growth approach of Jones (1995), Kortum (1997), and Segerstrom (1998).²⁴ These models propose that the idea production function takes the form

$$(17) \quad \frac{\dot{A}_t}{A_t} = (\alpha A_t^{-\beta}) \cdot S_t.$$

²⁴Jones (2005) provides an overview of this class of models. Atkeson and Burstein (2019); Atkeson, Burstein, and Chatzikonstantinou (2019); and Buera and Oberfield (2019) are recent examples.

Research productivity declines as A_t rises, so that it gets harder and harder to generate constant exponential growth. The elasticity β governs this process. That is, it parameterizes the extent to which ideas are getting harder to find.²⁵

Comparing both sides of equation (17), one can see that constant exponential growth requires a growing number of researchers S_t . In fact, if \dot{A}_t/A_t is constant over time, it must be that

$$(18) \quad g_A = \frac{g_S}{\beta},$$

where g_x denotes the constant growth rate of any variable x . The growth rate of the economy equals the growth rate of research effort deflated by the extent to which ideas are getting harder to find. Rising research effort and declining research productivity offset, endogenously in this framework, to deliver constant exponential growth.²⁶

This analysis can be applied across different firms, goods, or industries, following the insights of Ngai and Samaniego (2011), who develop a semi-endogenous growth model with heterogeneity in the dynamic spillover parameters of the idea production functions. Some goods, like semiconductors, can have rapid productivity growth because their β is small, while other goods like the speed of airplanes or perhaps the education industry itself could grow slowly because their β is large. Different sectors can exhibit constant exponential growth at different rates provided the amount of research effort put toward innovation is itself growing exponentially.

This framework helps us address a phenomenon that might at first have appeared puzzling: research productivity is declining very rapidly in the fastest growing sector in the economy, semiconductors. Why? In particular, why are we throwing so many resources at a sector that has such sharp declines in research productivity?

The last column of Table 7 reports estimates of β for each of our case studies, according to equation (17), and the results speak to the semiconductor puzzle we just highlighted. In particular, semiconductors is the application with the smallest value of β , coming in at 0.2, suggesting that it is the sector with the *least* degree of diminishing returns in idea production: A is growing at 35 percent per year, while research productivity is falling at 7 percent per year. From (18), this implies a value of β of $7/35 = 0.2$. In contrast, economy-wide TFP growth averages about 1.5 percent per year, while research productivity is declining at a rate of about 5 percent per year, yielding a β of more than 3! So in fact, semiconductors shows much *less* diminishing returns than the economy as a whole.

²⁵To map this structure into Jones (2005), notice Jones had $\dot{A} = \alpha S_t A_t^\phi$ so that $\beta = 1 - \phi$, whereas the Segerstrom and Kortum quality-ladder approaches are more naturally expressed in terms of β directly. The advantage of the β formulation is that it applies to both expanding variety and quality ladder models.

²⁶Many researchers in the micro-productivity literature work with an R&D-augmented production function in the spirit of Griliches (1994). This leads to a specification where a firm's own R&D (and other firms' spillover R&D) is an explanatory variable in an output-based production function (or other performance measures such as innovation proxies like patenting or market value as a forward-looking measure). Examples include Jaffe (1986); Hall and Mairesse (1995); Branstetter (2001); and Bloom, Schankerman, and Van Reenen (2013). Jones and Williams (1998) and Lucking, Bloom, and Van Reenen (2017) show how these Griliches-style production functions are compatible with semi-endogenous R&D growth models like those explored above, and therefore with the declining research productivity that is implied.

In the case where the “stepping on toes” effect of $\lambda \neq 1$ is allowed so that $\dot{A}_t/A_t = \alpha S_t^\lambda A_t^{-\beta}$, the steady-state equation becomes $g_A = \lambda g_S/\beta$, so that the estimate of β needs to be scaled by λ .²⁷ Appendix Table A1 reports the estimates of β when $\lambda = 0.75$.

Research productivity for semiconductors falls so rapidly, not because that sector has the sharpest diminishing returns; the opposite is true. It is instead because research in that sector is growing more rapidly than in any other part of the economy, pushing research productivity down. A plausible explanation for the rapid research growth in this sector is the “general purpose” nature of information technology. Demand for better computer chips is growing so fast that it is worth suffering the declines in research productivity there in order to achieve the gains associated with Moore’s Law.

B. Selection of Cases and Measurement Issues

Now that we have presented and discussed our evidence, it is worth stepping back to address general issues of selection and measurement. In particular, how did we pick our cases and does this create an important selection bias? Why do we focus on certain measures like density of semiconductors or seed yield per acre and not others?

As explained in the introduction, our selection of cases is driven primarily by the requirement that we are able to obtain data on both the “idea output” and the “research input” that is relevant for those ideas. We would like to report results for as many cases as possible, but this constraint was often binding. In the end, we report the cases in which we felt most confident. Another set of questions is why do we use the particular measures of idea output that we use. For example on Moore’s Law, why not look at the *price* of a floating point operation instead? Or why not output per farmer for crops instead of yield per acre? Or why not total factor productivity growth everywhere?

These are all good questions, and we have several answers. First and foremost, it is important to appreciate how high the hurdle is for issues like this to overturn our basic conclusion. Our graphs generally show flat or declining growth rates for idea output and research input measures that increase at 5 percent or 7 percent per year, or even more. That is, the research input doubles every 10 to 15 years. To overturn our result, alternative measures would need to dramatically change the pattern of growth rates we observe. Where our current measure is flat or declining, the alternative measure would need to show growth rates that themselves double every 10 to 15 years.

This seems extremely unlikely. In some cases, it can be easily checked. For example, Moore’s Law is sometimes stated in terms of cost as the cost of a floating point operation falling in half every two years. But this gives the same 35 percent growth rate that we’ve used and therefore delivers identical results to what we have. Next, notice that a *constant* mismeasurement of growth rates, i.e., if true growth rates of seed yields are 1 percent per year higher than we measure, similarly

²⁷ Our implementation is slightly richer, allowing for possible trends in \dot{A}_t/A_t itself. In particular, motivated by (17), we measure β as the ratio of the growth rate of research productivity to the growth rate of idea output.

leaves our main results essentially unchanged. The same would be true for measurement error of growth rates that is stationary. This is because our findings are driven by the enormous differences in trends in idea output and research input that we observe. What would change our results is if the measurement error was worsening systematically over time by a large amount. We know of no evidence suggesting this is the case.²⁸

A related point is that for our exercise to be valid, we need good matched measures of idea output and research input. For example, if the focus of US agricultural research is shifting to raising seed yields in other countries and climates, that could bias our results. Similarly, we would love to be able to discuss yield per farmer in addition to yield per acre. That would be interesting. The problem is that capital (e.g., new tractors and combines and GPS systems) has increasingly substituted for labor over time. This means that the R&D related to tractors and GPS systems then has first-order relevance. We think yield per acre is more tightly tied to the seed R&D that we are measuring.

The bottom line is that we've tried to be as careful as possible in measurement and that influences the cases we look at and the specific input and output measures we use.

Another point worth emphasizing is that we are measuring a "Solow residual" of the idea production function. Any inputs we have incorrectly omitted, such as business improvements from workers on the factory floor, the innovation efforts of failed start-ups that do not make it into the R&D numbers, or R&D from other industries that creates positive knowledge spillovers, will be forced into the residual. Our robust finding is that our measure of research productivity is declining. Missing inputs could affect our calculations in several ways. First, if the missing input enters as a spillover (outside the production function that researchers see), then it will naturally be part of our research productivity measure. Our conclusion, then, is that research productivity, including these spillover effects, is declining. Second, if the missing input belongs inside the research production function that firms see, the nature of the mismeasurement of research productivity depends on whether the input grows faster or slower than our measured research effort. If they grow at the same rate, research productivity will be correctly measured. A bias in the "bad" direction would occur if the unmeasured input grew more slowly than research effort. In this case, we would mistakenly overstate the decline in research productivity, when part should be attributed to another slower-growing input. Of course, it could just as easily be the case that the unmeasured input grows faster, in which case the bias works in our favor. We try to be careful in our measurement and hope that the robustness of our finding across many different settings mitigates these concerns.²⁹

Finally, we also face a trade-off in this paper: how much detail to go into in each particular case versus providing a sufficient number of cases so that the evidence

²⁸For example Aghion et al. (2017) finds relatively stable measurement error in growth rates.

²⁹Many authors have been concerned that economists are increasingly underestimating the growth of "intangible capital" inputs; e.g., see Corrado, Hulten, and Sichel (2005). One related concern is that entrant innovation efforts are not being fully counted, and this is becoming an increasing problem: think of the "garage" innovation in Silicon Valley. To the extent we increasingly fail to measure a rapidly-growing input, this will make the fall in research productivity even more dramatic.

is broadly convincing. An entire paper could be written on each of our cases. And we hope future researchers will dive into each of our cases to ensure that we have not missed something. We tried to strike a balance between depth and breadth of evidence, but different people make strike this balance differently. We hope that this project inspires others to look more closely at our individual cases and to delve into other cases where they can come up with good measures.

VIII. Conclusion

A key assumption of many endogenous growth models is that a constant number of researchers can generate constant exponential growth. We show that this assumption corresponds to the hypothesis that the total factor productivity of the idea production function is constant, and we proceed to measure research productivity in many different contexts.

Our robust finding is that research productivity is falling sharply everywhere we look. Taking the US aggregate number as representative, research productivity falls in half every 13 years: ideas are getting harder and harder to find. Put differently, just to sustain constant growth in GDP per person, the United States must double the amount of research effort every 13 years to offset the increased difficulty of finding new ideas.

This analysis has implications for the growth models that economists use in our own research, like those cited in the introduction. The standard approach in recent years employs models that assume constant research productivity, in part because it is convenient and in part because the earlier literature has been interpreted as being inconclusive on the extent to which this is problematic. We believe the empirical work we have presented speaks clearly against this assumption. A first-order fact of growth empirics is that research productivity is falling sharply.

Future work in the growth literature should determine how best to understand this fact. One possibility is the semi-endogenous growth models discussed in the preceding section. These models have important implications. For example, they have a “Red Queen” prediction in which we have to run faster and faster to maintain constant exponential growth.³⁰ If the growth rate of research inputs were to slow, this could cause economic growth itself to slow down. It is possible that this contributes to the global slowdown in productivity growth during the past 15 years.

Alternatively, there are other possible explanations for declining research productivity. Acigit and Kerr (2018) suggests that “follow on” innovations may be smaller than original innovations and provide evidence that research productivity declines with firm size. Incumbent firms may shift to “defensive” R&D to protect their market position, and this could cause research productivity to decline; Dinopoulos and Syropoulos (2007) provides a model along these lines. Or perhaps declines in basic research spending, potentially related to the US decline

³⁰Recall Lewis Carroll’s *Through the Looking-Glass*: “Now, here, you see, it takes all the running you can do, to keep in the same place. If you want to get somewhere else, you must run at least twice as fast as that!”

in publicly-funded research as a share of GDP, have negatively impacted overall research productivity. Clearly this would have important policy implications.³¹

That one particular aspect of endogenous growth theory should be reconsidered does not diminish the contribution of that literature. Quite the contrary. The only reason models with declining research productivity can sustain exponential growth in living standards is because of the key insight from that literature: ideas are nonrival. For example, if research productivity were constant, sustained growth would actually not require that ideas be nonrival; Akcigit, Celik, and Greenwood (2016) shows that rivalrous ideas can generate sustained exponential growth in this case. Our paper therefore suggests that a fundamental contribution of endogenous growth theory is not that research productivity is constant or that subsidies to research can necessarily raise growth. Rather it is that ideas are different from all other goods in that they can be used simultaneously by any number of people. Exponential growth in research leads to exponential growth in A_t . And because of nonrivalry, this leads to exponential growth in per capita income.³²

APPENDIX: ROBUSTNESS RESULTS: ALLOWING $\lambda < 1$

Table A1 shows the robustness of our results to the baseline assumption that there is no diminishing returns to research in the idea production function. See the discussion in Section IID for more details. For this set of results, we assume the input into the idea production function is S^λ where $\lambda = 3/4$. As expected, the growth rates of research productivity are about three-fourths as large as in the baseline case, reported in Table 7, but they are substantially negative nearly everywhere we look.

³¹ For example, see Akcigit, Hanley, and Serrano-Velarde (2016).

³² A similar point can be made with respect to the Schumpeterian models. For example, one reading of Aghion and Howitt (1992) is that it emphasizes both Schumpeterian creative destruction as well as that policy can affect long-run growth via an idea production function with constant research productivity. However, even if one accepts our interpretation of the evidence about research productivity, that in no way says anything critical about the Schumpeterian approach. Creative destruction and constant research productivity are distinct concepts.

TABLE A1—ROBUSTNESS RESULTS FOR RESEARCH PRODUCTIVITY: $\lambda = 0.75$

Scope	Time period	Average annual growth rate (%)	Half-life (years)	Dynamic diminishing returns, β
Aggregate economy	1930–2015	−4.0	17	2.4
Moore's Law	1971–2014	−5.1	14	0.15
Semiconductor TFP growth	1975–2011	−4.0	18	0.3
Agriculture, US R&D	1970–2007	−3.2	21	2.0
Agriculture, global R&D	1980–2010	−5.1	14	3.1
Corn, version 1	1969–2009	−7.9	9	5.7
Corn, version 2	1969–2009	−5.2	13	3.7
Soybeans, version 1	1969–2009	−5.4	13	4.6
Soybeans, version 2	1969–2009	−3.2	22	2.7
Cotton, version 1	1969–2009	−1.9	37	1.4
Cotton, version 2	1969–2009	+1.6	−44	−1.2
Wheat, version 1	1969–2009	−5.0	14	5.6
Wheat, version 2	1969–2009	−2.9	24	3.2
New molecular entities	1970–2015	−2.0	34	...
Cancer (all), publications	1975–2006	+0.4	−166	...
Cancer (all), trials	1975–2006	−3.4	20	...
Breast cancer, publications	1975–2006	−4.7	15	...
Breast cancer, trials	1975–2006	−7.7	9	...
Heart disease, publications	1968–2011	−2.8	25	...
Heart disease, trials	1968–2011	−5.4	13	...
Compustat, sales	3 decades	−9.4	7	0.9
Compustat, market cap	3 decades	−7.8	9	0.8
Compustat, employment	3 decades	−12.8	5	1.6
Compustat, sales/employment	3 decades	−3.9	18	0.9

Notes: This table shows results paralleling those in Table 7 when we allow for point-in-time diminishing returns to research. In particular, we assume the input into the idea production function is S^λ where $\lambda = 3/4$, as opposed to our baseline case of $\lambda = 1$.

REFERENCES

- Acemoglu, Daron, and Pascual Restrepo.** 2016. "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment." Unpublished.
- Acs, Zoltan J., and David B. Audretsch.** 1988. "Innovation in Large and Small Firms: An Empirical Analysis." *American Economic Review* 78 (4): 678–90.
- Aghion, Philippe, Ufuk Akcigit, and Peter Howitt.** 2014. "What Do We Learn from Schumpeterian Growth Theory?" In *Handbook of Economic Growth*, Vol. 2, edited by Philippe Aghion and Steven N. Durlauf, 515–63. Amsterdam: Elsevier.
- Aghion, Philippe, Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li.** 2017. "Missing Growth from Creative Destruction." NBER Working Paper 24023.
- Aghion, Philippe, and Peter Howitt.** 1992. "A Model of Growth through Creative Destruction." *Econometrica* 60 (2): 323–51.
- Aghion, Philippe, and Peter Howitt.** 1998. *Endogenous Growth Theory*. Cambridge, MA: MIT Press.
- Akcigit, Ufuk, Murat Alp Celik, and Jeremy Greenwood.** 2016. "Buy, Keep, or Sell: Economic Growth and the Market for Ideas." *Econometrica* 84 (3): 943–84.
- Akcigit, Ufuk, Douglas Hanley, and Nicolas Serrano-Velarde.** 2016. "Back to Basics: Basic Research Spillovers, Innovation Policy and Growth." Unpublished.
- Akcigit, Ufuk, and William R. Kerr.** 2018. "Growth through Heterogeneous Innovations." *Journal of Political Economy* 126 (4): 1374–1443.
- Akcigit, Ufuk, and Qingmin Liu.** 2016. "The Role of Information in Innovation and Competition." *Journal of the European Economic Association* 14 (4): 828–70.
- Atkeson, Andrew, and Ariel Burstein.** 2019. "Aggregate Implications of Innovation Policy." *Journal of Political Economy* 127 (6): 2625–83.
- Atkeson, Andrew, Ariel T. Burstein, and Manolis Chatzikonstantinou.** 2019. "Transitional Dynamics in Aggregate Models of Innovative Investment." *Annual Review of Economics* 11: 273–301.

- Bartelsman, Eric, and Wayne Gray.** 1996. "The NBER Manufacturing Productivity Database." NBER Working Paper 205.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta.** 2013. "Cross-Country Differences in Productivity: The Role of Allocation and Selection." *American Economic Review* 103 (1): 305–34.
- Beintema, Nienke M., Gert-Jan Stads, Keith O. Fuglie, and Paul Heisey.** 2012. *ASTI Global Assessment of Agricultural R&D Spending*. Washington, DC: International Food Policy Research Institute.
- Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb.** 2020. "Replication Data for: Are Ideas Getting Harder to Find?" American Economic Association [publisher], Inter-university Consortium for Political and Social Research [distributor]. <http://doi.org/10.3886/E111743V1>.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen.** 2013. "Identifying Technology Spillovers and Product Market Rivalry." *Econometrica* 81 (4): 1347–93.
- Blundell, Richard, Rachel Griffith, and John Van Reenen.** 1999. "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms." *Review of Economic Studies* 66 (3): 529–54.
- Branstetter, Lee G.** 2001. "Are Knowledge Spillovers International or Intranational in Scope?: Micro-econometric Evidence from the U.S. and Japan." *Journal of International Economics* 53 (1): 53–79.
- Buera, Francisco J., and Ezra Oberfield.** 2019. "The Global Diffusion of Ideas." Unpublished.
- Bureau of the Census.** 2009. "Survey of Industrial Research and Development 1972–2006." Federal Statistical Research Data Center.
- Bureau of the Census.** 2015a. "Business Research & Development and Innovation Survey 2007–2012." Federal Statistical Research Data Center.
- Bureau of the Census.** 2015b. "Census of Manufactures 1972–2012." Federal Statistical Research Data Center.
- Bureau of Economic Analysis.** 2017. "National Income and Product Accounts." Federal Reserve Economic Data (FRED) database. <https://fred.stlouisfed.org/> (accessed January 31, 2017).
- Bureau of Labor Statistics.** 2017. "Historical Multifactor Productivity Measures, 1948–2014." <https://www.bls.gov/mfp/home.htm> (accessed February 17, 2017).
- Centre for Industry Education Collaboration.** 2016. "The Essential Chemical Industry." <http://www.essentialchemicalindustry.org/> (accessed September 8, 2016).
- Clio Infra.** 2016. International Institute of Social History. <https://clio-infra.eu/> (accessed July 11, 2016).
- Compustat Industrial Annual Data.** 2016. Wharton Research Data Service. <https://wrds-www.wharton.upenn.edu/> (accessed June 20, 2016).
- Corrado, Carol, Charles Hulten, and Daniel Sichel.** 2005. "Measuring Capital and Technology: An Expanded Framework." In *Measuring Capital in the New Economy*, edited by Carol Corrado, John Haltiwanger, and Dan Sichel, 11–46. Chicago: University of Chicago Press.
- Cowen, Tyler.** 2011. *The Great Stagnation: How America Ate All the Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better*. New York: Dutton.
- Dalgaard, Carl-Johan, and Holger Strulik.** 2014. "Optimal Aging and Death: Understanding the Preston Curve." *Journal of the European Economic Association* 12 (3): 672–701.
- Davis, Steven J., John Haltiwanger, and Scott Schuh.** 1996. "Small Business and Job Creation: Dissecting the Myth and Reassessing the Facts." *Small Business Economics* 8 (4): 297–315.
- Dinopoulos, Elias, and Constantinos Syropoulos.** 2007. "Rent Protection as a Barrier to Innovation and Growth." *Economic Theory* 32 (2): 309–32.
- Dinopoulos, Elias, and Peter Thompson.** 1998. "Schumpeterian Growth without Scale Effects." *Journal of Economic Growth* 3 (4): 313–35.
- Evenson, Robert.** 1984. "International Invention: Implications for Technology Market Analysis." In *R&D, Patents and Productivity*, edited by Zvi Griliches, 89–126. Chicago: University of Chicago Press.
- Evenson, Robert E.** 1991. "Patent Data by Industry: Evidence for Invention Potential Exhaustion?" In *Technology and Productivity: The Challenges for Economic Policy*, edited by OECD, 233–48. Paris: OECD.
- Evenson, Robert E.** 1993. "Patents, R&D and Invention Potential: International Evidence." *American Economic Review* 83 (2): 463–68.
- Fernandez-Cornejo, Jorge, Jonathan Keller, David Spielman, Mohinder Gill, John King, and Paul Heisey.** 2004. *The Seed Industry in U.S. Agriculture: An Exploration of Data and Information on Crop Seed Markets, Regulation, Industry Structure, and Research and Development*. Washington, DC: United States Department of Agriculture, Economic Research Service.
- Flamm, Kenneth.** 2017. "Has Moore's Law Been Repealed? An Economist's Perspective." *Computing in Science and Engineering* 19 (2): 29–40.
- Fuglie, Keith, Paul Heisey, John King, Carl E. Pray, Kelly Day-Rubenstein, David Schimmelpfennig, Sun Ling Wang, and Rupa Karmarkar-Deshmukh.** 2011. "Research Investments and Market

- Structure in the Food Processing, Agricultural Input, and Biofuel Industries Worldwide.” *USDA-ERS Economic Research Report* (130).
- Gallini, Nancy T.** 2002. “The Economics of Patents: Lessons from Recent U.S. Patent Reform.” *Journal of Economic Perspectives* 16 (2): 131–54.
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow.** 2016. “How Destructive Is Innovation?” Unpublished.
- Gordon, Robert J.** 2016. *The Rise and Fall of American Growth: The US Standard of Living since the Civil War*. Princeton, NJ: Princeton University Press.
- Griliches, Zvi.** 1994. “Productivity, R&D and the Data Constraint.” *American Economic Review* 84 (1): 1–23.
- Grossman, Gene M., and Elhanan Helpman.** 1991. *Innovation and Growth in the Global Economy*. Cambridge, MA: MIT Press.
- Ha, Joonkyung, and Peter Howitt.** 2007. “Accounting for Trends in Productivity and R&D: A Schumpeterian Critique of Semi-Endogenous Growth Theory.” *Journal of Money, Credit and Banking* 39 (4): 733–74.
- Hall, Bronwyn H., and Jacques Mairesse.** 1995. “Exploring the Relationship between R&D and Productivity in French Manufacturing Firms.” *Journal of Econometrics* 65 (1): 263–93.
- Hall, Bronwyn H., Jacques Mairesse, and Pierre Mohnen.** 2010. “Measuring the Returns to R&D.” In *Handbook of the Economics of Innovation*, Vol. 2, edited by Bronwyn H. Hall and Nathan Rosenberg, 1033–82. Amsterdam: Elsevier.
- Hall, Robert E., and Charles I. Jones.** 2007. “The Value of Life and the Rise in Health Spending.” *Quarterly Journal of Economics* 122 (1): 39–72.
- Henry, Matthew D., and John L. Turner.** 2006. “The Court of Appeals for the Federal Circuit’s Impact on Patent Litigation.” *Journal of Legal Studies* 35 (1): 85–117.
- Hopenhayn, Hugo A.** 1992. “Entry, Exit, and Firm Dynamics in Long-Run Equilibrium.” *Econometrica* 60 (5): 1127–50.
- Howitt, Peter.** 1999. “Steady Endogenous Growth with Population and R&D Inputs Growing.” *Journal of Political Economy* 107 (4): 715–30.
- Hsieh, Chang-Tai, and Peter J. Klenow.** 2009. “Misallocation and Manufacturing TFP in China and India.” *Quarterly Journal of Economics* 124 (4): 1403–48.
- Human Mortality Database.** 2015. University of California, Berkeley. www.mortality.org (accessed November 18, 2015).
- Huffman, Wallace E., and Robert E. Evenson.** 2006. *Science for Agriculture: A Long-Term Perspective*. Hoboken, NJ: John Wiley & Sons.
- Jaffe, Adam B.** 1986. “Technological Opportunity and Spillovers of R&D: Evidence from Firms’ Patents, Profits and Market Value.” *American Economic Review* 76 (5): 984–1001.
- Jaffe, Adam B., and Josh Lerner.** 2006. “Innovation and Its Discontents.” *Capitalism and Society* 1 (3): Article 3.
- Jones, Benjamin F.** 2009. “The Burden of Knowledge and the ‘Death of the Renaissance Man’: Is Innovation Getting Harder?” *Review of Economic Studies* 76 (1): 283–317.
- Jones, Benjamin F.** 2010. “Age and Great Invention.” *Review of Economics and Statistics* 92 (1): 1–14.
- Jones, Charles I.** 1995. “R&D-Based Models of Economic Growth.” *Journal of Political Economy* 103 (4): 759–84.
- Jones, Charles I.** 2005. “Growth and Ideas.” In *Handbook of Economic Growth*, Vol. 1B, edited by Philippe Aghion and Steven A. Durlauf, 1063–1111. Oxford: North Holland.
- Jones, Charles I., and Jihee Kim.** 2018. “A Schumpeterian Model of Top Income Inequality.” *Journal of Political Economy* 126 (5): 1785–1826.
- Jones, Charles I., and John C. Williams.** 1998. “Measuring the Social Return to R&D.” *Quarterly Journal of Economics* 113 (4): 1119–35.
- Klette, Tor Jakob, and Samuel Kortum.** 2004. “Innovating Firms and Aggregate Innovation.” *Journal of Political Economy* 112 (5): 986–1018.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman.** 2015. “Technological Innovation, Resource Allocation, and Growth.” Unpublished.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman.** 2017. “Technological Innovation, Resource Allocation, and Growth.” *Quarterly Journal of Economics* 132 (2): 665–712.
- Kortum, Samuel.** 1993. “Equilibrium R&D and the Patent–R&D Ratio: U.S. Evidence.” *American Economic Review* 83 (2): 450–57.
- Kortum, Samuel.** 1997. “Research, Patenting, and Technological Change.” *Econometrica* 65 (6): 1389–1419.

- Kortum, Samuel, and Josh Lerner.** 1998. "Stronger Protection or Technological Revolution: What Is behind the Recent Surge in Patenting?" *Carnegie-Rochester Conference Series on Public Policy* 48: 247–304.
- Kremer, Michael.** 1993. "Population Growth and Technological Change: One Million B.C. to 1990." *Quarterly Journal of Economics* 108 (3): 681–716.
- Kruse-Andersen, Peter K.** 2016. "Testing R&D-Based Endogenous Growth Models." Unpublished.
- Laincz, Christopher A., and Pietro F. Peretto.** 2006. "Scale Effects in Endogenous Growth Theory: An Error of Aggregation Not Specification." *Journal of Economic Growth* 11 (3): 263–88.
- Lanjouw, Jean O., and Mark Schankerman.** 2004. "Patent Quality and Research Productivity: Measuring Innovation with Multiple Indicators." *Economic Journal* 114 (495): 441–65.
- Li, Chol-Won.** 2000. "Endogenous vs. Semi-Endogenous Growth in a Two-R&D-Sector Model." *Economic Journal* 110 (462): 109–22.
- Lichtenberg, Frank R.** 2017. "The Impact of Public and Private Research Support on Premature Cancer Mortality and Hospitalization in the U.S., 1999–2013." NBER Working Paper 23241.
- Lichtenberg, Frank R.** 2018. "The Long-Run Impact of New Medical Ideas on Cancer Survival and Mortality." NBER Working Paper 25328.
- Lucas, Robert E., Jr.** 1978. "On the Size Distribution of Business Firms." *Bell Journal of Economics* 9 (2): 508–23.
- Lucking, Brian, Nicholas Bloom, and John Van Reenen.** 2017. "Have R&D Spillovers Changed?" Unpublished.
- Melitz, Marc J.** 2003. "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity." *Econometrica* 71 (6): 1695–1725.
- Ngai, L. Rachel, and Roberto M. Samaniego.** 2011. "Accounting for Research and Productivity Growth across Industries." *Review of Economic Dynamics* 14 (3): 475–95.
- National Center for Health Statistics.** 2014. "Health, United States 2013." DHHS Publication 2014–1232.
- Nordhaus, William D.** 1997. "Do Real Output and Real Wage Measures Capture Reality? The History of Lighting Suggests Not." In *The Economics of New Goods*, edited by Timothy F. Bresnahan and Robert J. Gordon, 27–70. Chicago: University of Chicago Press.
- Nordhaus, William D.** 2003. "The Health of Nations: The Contribution of Improved Health to Living Standards." In *Measuring the Gains from Medical Research: An Economic Approach*, edited by Kevin M. Murphy and Robert Topel, 9–40. Chicago: University of Chicago Press.
- Oeppen, Jim, and James W. Vaupel.** 2002. "Broken Limits to Life Expectancy." *Science* 296 (5570): 1029–31.
- Pammolli, Fabio, Laura Magazzini, and Massimo Riccaboni.** 2011. "The Productivity Crisis in Pharmaceutical R&D." *Nature Reviews of Drug Discovery* 10 (6): 428–38.
- Pardey, Philip G., Connie Chan-Kang, Jason M. Beddow, and Steven P. Dehmer.** 2016. *Shifting Ground: Food and Agricultural R&D Spending Worldwide, 1960–2011*. St. Paul, MN: International Science and Technology Practice and Policy (InSTePP) Center at the University of Minnesota.
- Peretto, Pietro F.** 1998. "Technological Change and Population Growth." *Journal of Economic Growth* 3 (4): 283–311.
- Peretto, Pietro F.** 2016a. "A Note on the Second Linearity Critique." Unpublished.
- Peretto, Pietro F.** 2016b. "Robust Endogenous Growth." Unpublished.
- Perrin, Richard K., K. A. Kunnings, and L. A. Ihnen.** 1983. *Some Effects of the U.S. Plant Variety Protection Act of 1970*. North Carolina State University Department of Economics and Business. Economics Research Report 46.
- Pillai, Unni.** 2016. "Vertical Specialization and Industry Growth: The Case of the Semiconductor Industry." Unpublished.
- Romer, Paul M.** 1990. "Endogenous Technological Change." *Journal of Political Economy* 98 (5): S71–102.
- Scannell, Jack W., Alex Blanckley, Helen Boldon, and Brian Warrington.** 2012. "Diagnosing the Decline in Pharmaceutical R&D Efficiency." *Nature Reviews Drug Discovery* 11 (3): 191–200.
- Segerstrom, Paul S.** 1998. "Endogenous Growth without Scale Effects." *American Economic Review* 88 (5): 1290–1310.
- Solow, Robert M.** 1957. "Technical Change and the Aggregate Production Function." *Review of Economics and Statistics* 39 (3): 312–20.
- Traxler, Greg, Albert K. A. Acquaye, Kenneth Frey, and Ann Marie Thro.** 2005. *Public Sector Plant Breeding Resources in the US: Study Results for the Year 2001*. Washington, DC: United States Department of Agriculture Cooperative State Research, Education and Extension Service.
- University of York.** 2016. "The Essential Chemical Industry." <http://www.essentialchemicalindustry.org/> (accessed September 8, 2016).

- US Department of Agriculture Economic Research Service.** 2018a. "Agricultural Productivity in the U.S." Table 1. <https://www.ers.usda.gov/data-products/agricultural-productivity-in-the-us.aspx> (accessed January 31, 2018).
- US Department of Agriculture Economic Research Service.** 2018b. "Agricultural Research Funding in the Public and Private Sectors, 1970–2009." <https://www.ers.usda.gov/data-products/agricultural-research-funding-in-the-public-and-private-sectors/> (accessed January 31, 2018).
- US Department of Agriculture National Agricultural Statistics Service.** 2016. <https://quickstats.nass.usda.gov> (accessed September 8, 2016).
- US Department of Agriculture National Institute of Food and Agriculture Current Research Information System.** 2016. <http://cris.nifa.usda.gov/fsummaries.html> (accessed September 8, 2016).
- Vaupel, James W., and Vladimir Canudas Romo.** 2003. "Decomposing Change in Life Expectancy: A Bouquet of Formulas in Honor of Nathan Keyfitz's 90th Birthday." *Demography* 40 (2): 201–16.
- Webb, Michael, Nick Short, Nicholas Bloom, and Josh Lerner.** 2018. "Some Facts of High-Tech Patenting." NBER Working Paper 24793.
- Wikipedia.** 2017. "Moore's Law." https://web.archive.org/web/20170126022932/http://en.wikipedia.org/wiki/Moore's_law (accessed January 3, 2017).
- Young, Alwyn.** 1998. "Growth without Scale Effects." *Journal of Political Economy* 106 (1): 41–63.
- Zeira, Joseph.** 2011. "Innovations, Patent Races and Endogenous Growth." *Journal of Economic Growth* 16 (2): 135–56.