

Are Ideas Getting Harder to Find?

Nicholas Bloom

Stanford University and NBER

Charles I. Jones

Stanford University and NBER

John Van Reenen

MIT and NBER

Michael Webb*

Stanford University

March 5, 2018 — Version 2.0

Abstract

In many models, economic growth arises from people creating ideas, and the long-run growth rate is the product of two terms: the effective number of researchers and their research productivity. We present a wide range of 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 every two years of the density of computer chips is more than 18 times larger than the number required in the early 1970s. Across a broad range of case studies at various levels of (dis)aggregation, we find that ideas — and in particular the exponential growth they imply — are getting harder and harder to find. Exponential growth results from large increases in research effort that offset its declining productivity.

*We are grateful to 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 the CEPR Macroeconomics and Growth Conference, George Mason, Harvard, the Minneapolis Fed, the NBER growth meeting, the NBER Macro across Time and Space conference, the Rimini Conference on Economics and Finance, and Stanford for helpful comments and to Antoine Dechezlepretre, Dietmar Harhoff, Wallace Huffman, Keith Fuglie, Unni Pillai, and Greg Traxler for extensive assistance with data. The Alfred Sloan Foundation, the European Research Council and ESRC have provided financial support. An online data appendix with replication files can be obtained from <http://web.stanford.edu/people/chadj>. See <http://web.stanford.edu/people/chadj/papers.html> for a History of Revisions to see what has changed.

1. Introduction

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:

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

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 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 around 35% 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 and harder to find. Research productivity in this case is declining sharply, at a rate that averages about 6.8% per year.

We document qualitatively similar results throughout the U.S. economy. We consider detailed microeconomic evidence on idea production functions, focusing on places where we can get the best measures of both the output of ideas and the inputs used to produce them. 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% per year. We find a similar rate of decline when studying the mortality improvements associated with cancer and heart disease.

Finally, we examine firm-level data from Compustat to provide another perspective. While the data quality from this sample is not as good as for our case studies, the latter suffer from possibly not being representative. We find substantial heterogeneity across firms, but research productivity is declining in more than 85% of our sample. Averaging across firms, research productivity declines at a rate of around 10% per year.

Perhaps research productivity is declining sharply within every particular case that we look at 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 for this reason, we present this application first to illustrate our methodology. We find that research productivity for the aggregate U.S. economy has declined by a factor of 41 since the 1930s, an average decrease of more than 5% per year.

This is a good place to explain why we think looking at the macro data is insufficient and why studying the idea production function at the micro level is crucial; Section 3 below discusses this issue in more detail. The overwhelming majority of papers on economic growth published in the past decade are based on models in 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

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.

²Examples are cited below after equation (1).

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

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 U.S. 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. Moreover, we suggest at the end of the paper that the rapid declines in research productivity that we see in semiconductors, for example, might be precisely due to the fact that research effort is rising so sharply. Because it gets harder to find new ideas as research progresses, a sustained and massive expansion of research like we see in semiconductors (for example, because of the “general purpose technology” nature of information technology) may lead to a substantial downward trend in research productivity.

Others have also provided evidence suggesting that ideas may be getting harder to find over time. Griliches (1994) provides a summary of the earlier literature exploring the decline in patents per dollar of research spending. Gordon (2016) reports extensive new historical evidence from throughout the 19th and 20th centuries. Cowen (2011) synthesizes earlier work to explicitly make the case. (Ben) Jones (2009) 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.

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 all of them, however, 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. However, 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 we felt most confident.

The remainder of the paper is organized as follows. Section 2 lays out our conceptual framework and presents the aggregate evidence on research productivity to illustrate our methodology. Section 3 places this framework in the context of growth theory and suggests that applying the framework to micro data is crucial for understanding the nature of economic growth. Sections 4 through 7 consider our applications to Moore’s Law, agricultural yields, medical technologies, and Compustat firms. Section 8 then revisits the implications of our findings for growth theory, and Section 9 concludes.

2. Research Productivity and Aggregate Evidence

2.1. The Conceptual Framework

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

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

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 (2016b), 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 20th 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:

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

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 sub-

sidies produce “permanent growth effects” are equivalent statements.⁴ 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.

2.2. 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 the main way that growth proceeds is by developing ever-better 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. Maybe research productivity is declining within every product line, but still not declining for the economy as a whole because we keep inventing new products. An analysis focused on microeconomic case studies would 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) model itself, and to generate constant research productivity in that case requires the equation we started the paper with:

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

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

⁴The 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.”

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 U.S. 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 4 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.⁵ But reasonable variations along these lines would not change the

⁵Figure 1 in the online appendix 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 US) 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

Figure 1: Aggregate Data on Growth and Research Effort




Note: 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 BLS 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 Robert Gordon (2016). The idea input measure is gross domestic investment in intellectual property products from the National Income and Product Accounts, deflated by a measure of the nominal wage for high-skilled workers.

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

3. Refining the Conceptual Framework

In this section, we further develop the conceptual framework introduced in the previous section. First, we explain why the aggregate evidence just presented can be misleading, motivating our focus on micro data. 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 we use.

Figure 2: Aggregate Evidence on Research Productivity



Note: Research productivity is the ratio of idea output, measured as TFP growth, to research effort. See notes to Figure 1 and the online data appendix. Both research productivity and research effort are normalized to the value of 1 in the 1930s.

sions to our calculation of research productivity.

3.1. The Importance of Micro Data

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*.⁶

⁶This 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; 2016b).

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,

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

Here, S_{it} is the number of scientists devoted to improving the quality of good i , and in a symmetric case, we might have $S_{it} = \frac{S_t}{N_t}$. The key is that the aggregate number of scientists S_t can be growing, but perhaps the number 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:

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

3.2. “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 ideas is researchers, then deflating R&D expenditures by an average wage will appropriately recover the effective quantity of researchers. In practice, R&D expenditures also typically involve 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 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

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

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:

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

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

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

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. First, notice that dividing both sides of the production function for final output by Y_t^θ and rearranging yields

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

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

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

And rearranging gives

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

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,⁷ (6) can be written as

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

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 $\tilde{S}_t := \frac{R_t}{w_t} = \frac{R_t}{Y_t} \cdot \frac{Y_t}{w_t} = s_t L / \bar{\theta}$. 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. Of course, 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

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

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:

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

The only difference is that we deflate R&D expenditures by a measure of the nominal

⁷The “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.

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) above 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, this naive approach to research productivity would not really correspond to anything economically useful.

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.⁸

3.3. 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

⁸These data are from Census Tables P18 and P19, available at <http://www.census.gov/topics/income-poverty/income/data/tables.html>. Prior to 1991, we use the series for “4 or more years of college.” For years between 1939 and 1967, we use the series Bc845 from the Historical Statistics for the U.S. Economy, Millennial Edition. Finally, for the aggregate research productivity calculation in Figure 1, we require a deflator from the 1930s. We extrapolate the college earnings series backward into the 1930s using nominal GDP per person for this purpose. 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 rise in research productivity. As an alternative, we redid all our results 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 similar, in part because the increases in research productivity that we document are so large.

effective number of researchers working to improve that product.⁹ 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.¹⁰ The two previous subsections then merge together very naturally.

3.4. Stepping on Toes

One other potential modification to the idea production function that has been considered in the literature is a duplication externality. For example, perhaps the idea production function depends on S_t^λ , where λ is less than one. 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 choose instead to focus on the benchmark case of $\lambda = 1$ for several reasons. First, our measurement 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 —

⁹The 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.

¹⁰Building 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.

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, Appendix A 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.

4. 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 fifty 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.¹¹

A doubling time of two years is equivalent to a constant exponential growth rate of 35 percent per year. While there is some discussion of Moore's Law slowing down in recent years (there always seems to be such discussion!), we will take the constant exponential growth rate as corresponding to a constant flow of new ideas back to 1971. That is, we assume the output of the idea production for Moore's Law is 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

¹¹See the Wikipedia entry at https://en.wikipedia.org/wiki/Moores_law.

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

Figure 3: The Steady Exponential Growth of Moore's Law




Source: Wikipedia, <https://en.wikipedia.org/wiki/Moores.law>.

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 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 below and in the online data 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

Figure 4: Data on Moore's Law



Note: 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 data includes 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.

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 NBER/CES Manufacturing Industry Database. Because of an acceleration in average TFP growth there from around

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
(a) Narrow	8	-4.8%	14.5
(b) Narrow (downweight conglomerates)	11	-5.6%	12.3
(c) Broad (downweight conglomerates)	26	-7.6%	9.1
(d) Intel only (narrow)	347	-13.6%	5.1
(e) Intel+AMD (narrow)	352	-13.6%	5.1
<i>TFP Growth in Semiconductors (1975–2011)</i>			
(f) Narrow (no equipment R&D)	5	-3.2%	21.4
(g) Narrow (with equipment R&D)	7	-4.4%	15.8
(h) Broad (no equipment R&D)	11	-5.6%	12.3
(i) Broad (with equipment R&D)	13	-6.1%	11.3

Note: The R&D measures are based on Compustat 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 U.S. 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% of Intel's patents were for semiconductors, but by 2010 this number had fallen to just 8%. 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 (d) and (e) show results when we consider the “Narrow” measure of R&D but focus on only one or two firms. Rows (f) through (i) 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 an 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 data appendix for more details.

8% per year to around 20% per year in the late 1990s — approximately a 3-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.

4.1. 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 science may lead to a new idea that improves computer chips. Such positive spillovers are not a problem for our analysis since they would show up as an *increase* in research productivity rather than as the declines that we document in this paper.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.

5. Agricultural Crop Yields

Our next application for measuring research productivity is agriculture. Due partly to the historical importance of agriculture in the economy, data on 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”

¹³If such spillovers were larger at the start of our time period than at the end, we could be underestimating the impact of semiconductor R&D on productivity growth. We do not know of evidence suggesting this. For example, see Lucking, Bloom and Van Reenen (2017).

Figure 5: TFP Growth and Research Effort in Agriculture

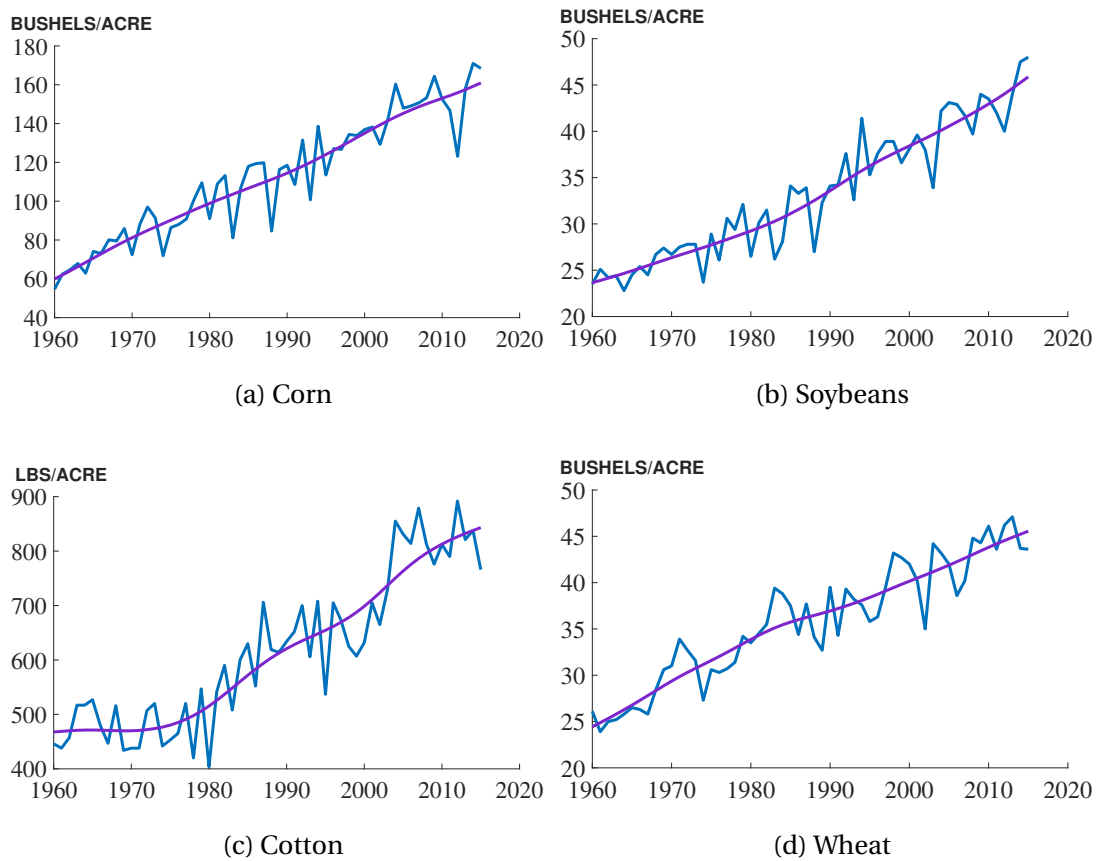


Note: Both TFP growth and U.S. R&D spending (public and private) for the agriculture sector as a whole are taken from the U.S. Department of Agriculture's Economic Research Service. The TFP series is smoothed with an HP filter. Global R&D spending for agriculture is taken from Fuglie, Heisey, King, Day-Rubenstein, Schimmelpfennig, Wang, Pray and Karmarkar-Deshmukh (2011), Beintema, Stads, Fuglie and Heisey (2012), and Pardey, Chan-Kang, Beddow and Dehmer (2016). Nominal R&D spending is deflated by the wage for college graduates, as described earlier.

measure, we use (a smoothed 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, the growth rate of yield per acre and total factor productivity growth look similar, and Figure 5 gives us no reason to believe that there is any large acceleration in total factor productivity growth in this sector.

Figure 6: U.S. Crop Yields



Note: Smoothed yields are computed using an HP filter with a smoothing parameter of 400.

For each of corn, soybeans, cotton, and wheat, we measure research inputs as R&D expenditure directed at improving crop yields. R&D is generally broken down into research on biological efficiency (cross-breeding and bioengineering), mechanization, management, protection and maintenance, and post-harvest (see, for example, Huffman and Evenson (2006)). We count research on biological efficiency and protection and maintenance as the portion devoted to improving crop yields.

Figure 6 shows yields for our four crops back to the 1960s, measured in bushels or pounds harvested per acre planted.¹⁴ These correspond to average yields realized on

¹⁴ For our measure of ideas, we use the national realized yields series for each crop available from the U.S. Department of Agriculture National Agricultural Statistics Service (2016). Our measure of R&D inputs consists of the sum of R&D spending by the public and private sectors in the U.S. 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 et al. (1983), Fernandez-

U.S. farms. They are therefore subject to many influences, including choice of inputs and random shocks. These shocks, especially adverse weather and pest events, tend to have asymmetric effects: adverse events cause much larger reductions in yields than favorable events increase them, as indicated by the many large one-year reductions followed by recoveries in the figure (see Huffman, Jin and Xu (2018)). Nevertheless, yields across these four crops roughly doubled between 1960 and 2015.

Figure 7 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.

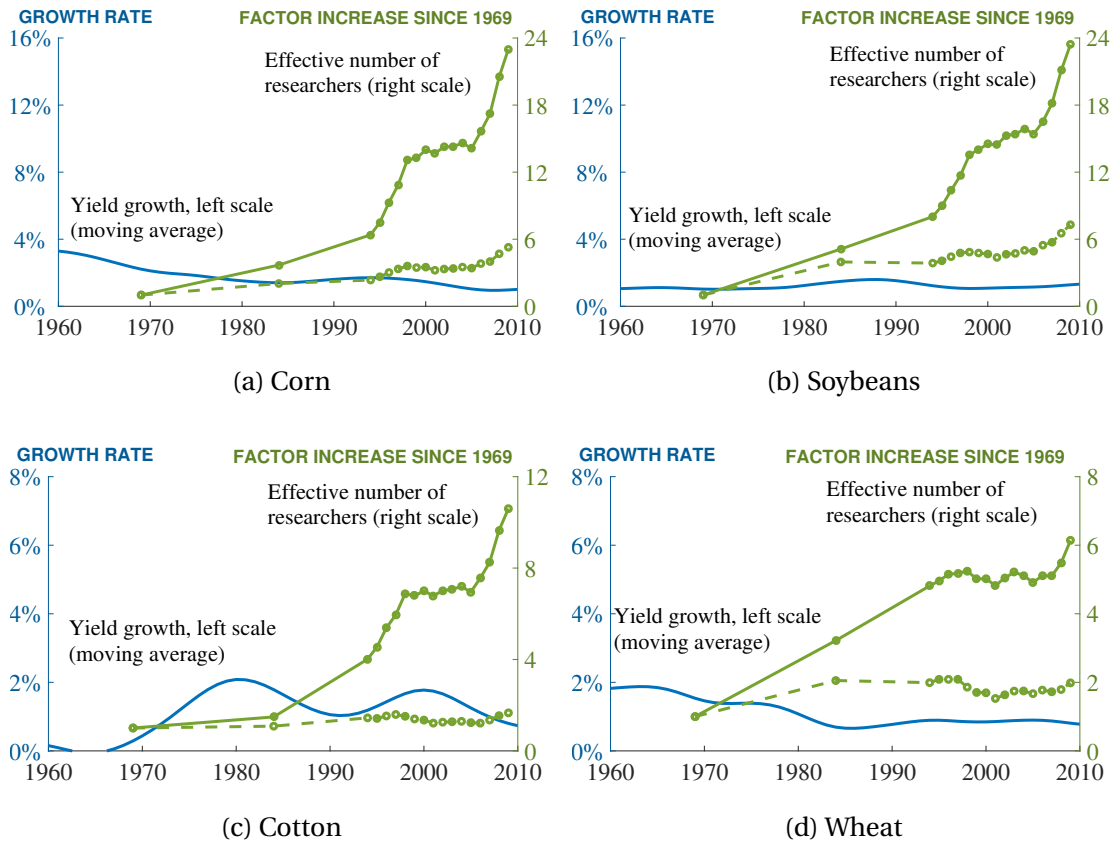
The green lines in Figure 7 show measures of the “effective” number of researchers working on each crop, measured as the sum of public and private R&D spending deflated by the wage of high-skilled workers. Two measures are presented. The faster-rising number corresponds to research targeted only at so-called 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 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.

It is immediately evident from Figure 7 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

Cornejo et al. (2004), Traxler et al. (2005), Huffman and Evenson (2006), and University of York (2016). Data on U.S. public sector R&D expenditure by crop are from the U.S. 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).

Figure 7: Yield Growth and Research Effort by Crop



Note: The blue line is the annual growth rate of the smoothed yields over the following 5 years, from Figure 6. The two green lines report “Effective Research”: the solid line is based on R&D targeting seed efficiency only; the dashed line additionally includes research on crop protection. Both are normalized to one in 1969. R&D expenditures are deflated by a measure of the nominal wage for high-skilled workers. See footnote 14 and the online data appendix for more details.

Table 2: Research Productivity in Agriculture, 1969–2009

Crop	— Effective research —		Research Productivity	
	Factor increase	Average growth	Factor decrease	Average growth
<i>Research on seed efficiency only</i>				
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%
<i>Research includes crop protection</i>				
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%
<i>Agriculture</i>				
U.S. research, 1970–2008	1.9	1.8%	3.9	-3.7%
Global research, 1980–2010	1.6	1.6%	5.2	-5.5%

Note: 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. R&D expenditures are deflated by a measure of the nominal wage for high-skilled workers. See footnote 14 and the online data appendix for more details.

23 for both corn and soybeans if we restrict attention to seed-yield research narrowly-defined. If yield growth were constant (which is not a bad approximation across the four crops as shown in Figure 7), then research productivity would on average 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.

6. Mortality and Life Expectancy


Health expenditures account for around 18 percent of U.S. 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.

6.1. New Molecular Entities

Our first example from the medical sector is a fact that is well-known in the literature, recast in terms of research productivity. Figure 8 shows the number of new molecular entities (“NMEs”) approved by the Food and Drug Administration. NMEs are new drugs, including both chemical and biological products, that have been approved by the FDA. Virtually all pharmaceutical advances in the last 50 years show up in these counts (Zambrowicz and Sands, 2003). Famous examples that became commercial blockbuster drugs are Zocor (for cholesterol), Prilosec (for gastroesophageal reflux), Claritin (for allergies), Celebrex (for arthritis), and Taxol (for treating various types of cancer). Only two or three of the NMEs in any given year become commercial successes. Among famous drugs, only morphine and aspirin do not show up in these counts, because their discovery pre-dates the FDA. The flow of NMEs is well-known to show very little trend, although 2014 and 2015 are two of the years with the most approvals. Based on this fact, we proceed conservatively and measure idea output as the flow of NMEs rather than as the percentage change.

We obtain data on pharmaceutical R&D spending from the Pharmaceutical Research and Manufacturers of America (Phrma), which has conducted an annual survey of its members back to 1970 and includes R&D performed both domestically and abroad

Figure 8: New Molecular Entities Approved by the FDA



Note: Historical data on NME approvals are from Food and Administration (2013). Data for recent years are taken from Pharmaceutical Research and Manufacturers of America (2016).


by these companies.¹⁵ Using the procedures described earlier, we get the research productivity and effective research numbers shown in Figure 9. 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.¹⁶

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, the NMEs still suffer from an important aggregation issue, adding up across a wide range of health conditions. These limitations motivate the approach described next, in which we focus on the productivity of medical research in specific diseases.

¹⁵A limitation is that it does not include R&D done by foreign companies that is performed abroad. However, Figure 1 of Congressional Budget Office (2006) suggests that this is still a very useful measure.

¹⁶It is well documented that the number of NMEs per dollar of R&D is declining; our statement is different in that, importantly, our measure of research input is deflated by the wage of college graduates. Akcigit and Liu (2016) examine this case in more detail and suggest that the rising replication of “dead-ends” in pharmaceutical research (and elsewhere) could be part of the story.

Figure 9: Research Productivity for New Molecular Entities




Note: Historical data on NME approvals are from Food and Administration (2013). Data on research spending by the pharmaceutical industry are from the 2010, 2013, and 2016 editions of Pharmaceutical Research and Manufacturers of America (2016). See the online data appendix for more details.

6.2. Years of Life Saved

To measure idea output in treating diseases, we begin with life expectancy. Figure 10 shows U.S. 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 provides a better description of the time path of life expectancy. Since 1950, U.S. 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.

Also shown in the graph is the well-known fact that overall life expectancy grew even more rapidly during the first half of the 20th 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

Figure 10: U.S. Life Expectancy



Source: *Health, United States 2013* and <https://www.clio-infra.eu>.

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 will 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.

6.3. 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

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

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

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

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.

Vaupel and Canudas-Romo (2003) show 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

As discussed at the start of this section, life expectancy tends to rise linearly. Therefore, constant exponential growth in income per person is associated with constant arithmetic increases in life expectancy, which in the aggregate average 1.8 years per decade in the U.S. We therefore take the quantity described in equation (16) as our measure of the output of ideas associated with declines in mortality for a given disease.¹⁸

The research input aimed at reducing mortality from a given disease is at first bluish 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

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.

¹⁸Our measures of life expectancy and mortality from all sources by age come from the Human Mortality Database at <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 at <http://seer.cancer.gov/>.

have “Neoplasms,” for example, as a MESH (Medical Subject Heading) term. MESH is the National Library of Medicine’s controlled vocabulary thesaurus.¹⁹ 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.²⁰

Figure 11 shows our basic “idea output” 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 1000 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 1000 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.


Figure 12 shows 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 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


¹⁹For 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/>.

²⁰Lichtenberg (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.


Figure 11: Mortality and Years of Life Saved



(a) All cancers




(b) Breast cancer




(c) Heart Disease

Note: For the two cancer figures, 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. The “Years of life saved per 1000 people” is computed using equation (16), as described in the text.


Figure 12: Medical Research Effort



(a) All cancers



(b) Breast cancer



(c) Heart Disease

Note: The number of publications and clinical trials are 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.

Table 3: Research Productivity for Medical Research

Disease	— Effective research —		Research Productivity	
	Factor increase	Average growth	Factor decrease	Average growth
<i>All publications</i>				
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%
<i>Clinical trials only</i>				
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%


Note: 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 data appendix for more details.

ratio of years of life saved to the number of publications. Figure 13 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.


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 one year by 2006.

Next, however, notice that the changes were not monotonic if we go back to 1975.


Figure 13: Research Productivity for Medical Research



(a) All cancers



(b) Breast cancer



(c) Heart Disease

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

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.

7. Research Productivity in Firm-Level Data

Our studies of semiconductors, crops, and medicine 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.²¹

To help address these concerns, we turn to Compustat data on US publicly-traded firms. The strength of these data is that they are more representative than the case studies, but of course they too have limitations. Publicly-traded firms are still a select sample, and our measures of “ideas” and research inputs are likely less precise. However, as a complement to the case studies, we find this evidence helpful.

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 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, Papanikolaou, Seru and Stoffman 2015). Patents are also positively correlated with the firm’s subsequent growth in employment and sales.

²¹For example, Peretto (1998) and Peretto (2016b) emphasize this perspective on varieties, while Aghion and Howitt (1992) take the alternative view that different firms may be involved in producing the same variety. Klette and Kortum (2004) allow the number of varieties produced by each firm to be heterogeneous and to evolve over time.

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.²² 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, 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 non-zero 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

1. 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.
2. 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).
3. 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,

²²This 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.

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.

Table 4 shows our research productivity calculation for various cuts of the Compustat data and following firms that we observe for two, three, and four decades. In all samples, there is substantial growth in the effective number of researchers within each firm, with growth rates averaging between 2.4% and 8.8% 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 steady, rapid declines in firm-level research productivity across all samples, at growth rates that range from -4.2% to -14.5% per year for multiple decades.

Averaging across all our samples, research productivity falls at a rate of about 9% 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.

The next three figures characterize the heterogeneity across firms in our Compustat sample by showing the distribution of the factor changes in effective research and research productivity across all the firms; to keep things manageable, we focus on the results for sales revenue, but the results with other output measures are similar. Figure 14 shows this distribution for the firms we observe for only two decades, while Figures 15 and 16 shows the distributions for the firms observed for three and four decades.

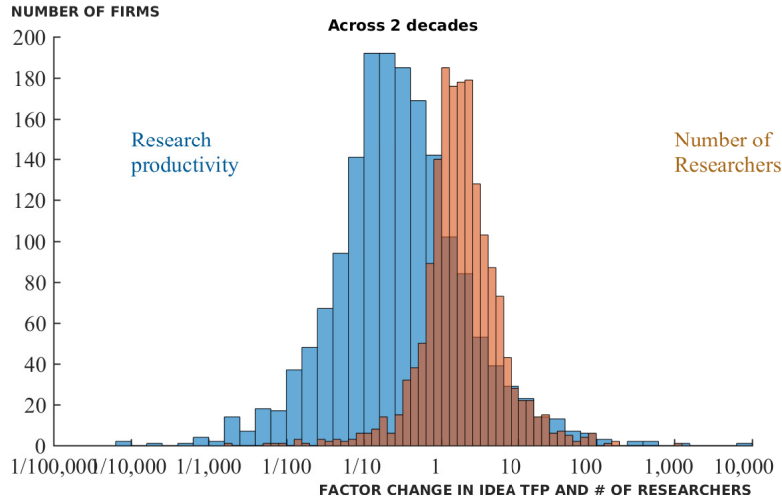
The heterogeneity across firms is impressive and somewhat reminiscent of the heterogeneity we see in our case studies. Nevertheless, it is clear from these histograms that there is essentially no evidence that constant research productivity is a good characterization of the firm-level data. The average, 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% of firms in any of the histograms have research

Table 4: Research Productivity in Compustat Firm-Level Data

Sample	— Effective research —		Research Productivity	
	Factor increase	Average growth	Factor decrease	Average growth
<i>Sales Revenue</i>				
2 decades (1712 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%
<i>Market Cap</i>				
2 decades (1124 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%
<i>Employment</i>				
2 decades (1395 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%
<i>Revenue Labor Productivity</i>				
2 decades (1444 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%


Note: 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. R&D expenditures are deflated by a measure of the nominal wage for high-skilled workers. The average growth rate across 2 decades is computed by dividing by 10 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 data appendix for more details.

Figure 14: Compustat Distributions, Sales Revenue (2 Decades)




Note: Based on 1712 firms. 22.1% of firms have increasing research productivity. Only 3.0% have research productivity that is roughly constant, defined as a growth rate whose absolute value is less than 1% per year.

Figure 15: Compustat Distributions, Sales Revenue (3 Decades)



Note: Based on 469 firms. 11.9% of firms have increasing research productivity. Only 4.3% have research productivity that is roughly constant, defined as a growth rate whose absolute value is less than 1% per year.

Figure 16: Compustat Distributions, Sales Revenue (4 Decades)



Note: Based on 149 firms. 14.8% of firms have increasing research productivity. 4.7% firms in this sample have research productivity that is roughly constant, defined as a growth rate whose absolute value is less than 1% per year.

productivity changing (either rising or falling) by less than 1% per year on average.

Table 5 provides additional evidence of the robustness of these results. In the interests of brevity, we report these results for the sales revenue output measure for firms that we observe across three decades, but the results for our other output measures and time frames are similar. The first row of the table repeats the benchmark results described earlier. The second and third rows relax the requirement that sales growth is positive. In the second row, we replace all growth rates less than 1% with 1%, while in the third row we additionally winsorize an equal number of firms at the top of the growth rate distribution. 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 in each decade rather than mean sales growth as our output measure. And the last row reports unweighted averages rather than weighting firms by the effective number of researchers. The general finding of substantial declines in research productivity is robust.

Table 5: Compustat Sales Data across 3 Decades: Robustness

Case	— Effective research —		Research Productivity	
	Factor increase	Average growth	Factor decrease	Average growth
Benchmark (469 firms)	3.8	6.7%	9.2	-11.1%
Winsorize $g < .01$ (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 <i>any</i> 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%

Note: Robustness results reported for the sample of changes across 3 decades. “Winsorize $g < .01$ ” means we replace any idea output measure that is less than 1% annually with a value of 1%. “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.

8. Discussion

8.1. Summary and Semi-Endogenous Growth

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 6, research productivity is declining at a substantial rate in virtually every place we look. The table also provides a way to 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 or not 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).²³ These models propose that the idea production function takes the form

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

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 measures the extent of dynamic diminishing returns in idea production, with a higher β meaning that research productivity declines more rapidly as A_t increases.

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

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

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 of diminishing returns in idea production. Rising research effort and declining research productivity offset — endogenously in this framework — to deliver constant exponential growth.

Finally, 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 have slow rates of innovation because their β is large. Nevertheless, all products could exhibit constant exponential growth if 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

²³Jones (2005) provides a broad overview of this class of models.

Table 6: Summary of the Evidence on Research Productivity

Scope	Time Period	Average annual growth rate	Half-life (years)	Extent of Diminishing Returns, β
Aggregate economy	1930–2015	-5.3%	13	3.4
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/emp	3 decades	-4.5%	15	1.1

Note: 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 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).

resources at a sector that has such sharp declines in research productivity?

The last column of Table 6 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% per year, while research productivity is falling at 7% per year. From (17), this implies a value of β of $7/35 = 0.2$. In contrast, economy-wide TFP growth averages about 1.5% per year, while research productivity is declining at a rate of about 5% per year, yielding a β of more than 3! So in fact, semiconductors shows much *less* diminishing returns than the economy as a whole.

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.

8.2. 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% or 7% 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 doubling 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 2 years. But this gives the same 35% 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% per year higher than we measure — similarly 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.²⁴

A related point is that for our exercise to be valid, we need good measures of idea output and research input. 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.

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

²⁴For example Aghion, Bergeaud, Boppart, Klenow and Li (2017) find relatively stable measurement error in growth rates.

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.²⁵

9. 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 U.S. 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 U.S. must double the amount of research effort searching for new ideas 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

²⁵Many 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 Reenen (2013). Jones and Williams (1998) show how these Griliches-style production functions are compatible with R&D growth models like those explored above — and therefore with the declining research productivity that is implied.

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.²⁶ 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 fifteen years.

One way in which the overall interpretation of the evidence in this paper could be wrong is if there is indeed declining research productivity in every sector, but the entire increase in aggregate R&D occurs in making quality/productivity improvements for individual varieties. The separate idea production function for producing new varieties could then exhibit constant research productivity if in fact the research going to create new varieties is not increasing over time. While this is conceptually possible, it seems unlikely. We do not have any growth models in the literature that suggest this could occur in an equilibrium in which there is population growth. Why, despite a growing population and constant research productivity, would the equilibrium allocation feature a constant number of researchers creating new varieties? Nevertheless, this loophole is one that future research may wish to consider.

Alternatively, there are other possible explanations for declining research productivity. Akcigit and Kerr (2018) suggest 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) provide models along these lines. Or perhaps declines in basic research spending (potentially related to the U.S. decline in publicly-funded research as a share of GDP) have negatively impacted overall research productivity. Clearly this would have important policy implications.²⁷

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

²⁶Recall 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!”

²⁷For example, see Akcigit, Hanley and Serrano-Velarde (2016a).

not require that ideas be nonrival; Akcigit, Celik and Greenwood (2016b) show that fully rivalrous ideas in a model with perfect competition 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 do not get depleted when used by more and more 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.

References

- Acemoglu, Daron and Pascual Restrepo, “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares and Employment,” May 2016. unpublished.
- Aghion, Philippe and Peter Howitt, “A Model of Growth through Creative Destruction,” *Econometrica*, March 1992, 60 (2), 323–351.
- and —, *Endogenous Growth Theory*, Cambridge, MA: MIT Press, 1998.
- , Antonin Bergeaud, Timo Boppart, Peter J. Klenow, and Huiyu Li, “Missing Growth from Creative Destruction,” NBER Working Papers 24023, National Bureau of Economic Research, Inc November 2017.
- , Ufuk Akcigit, and Peter Howitt, “What Do We Learn From Schumpeterian Growth Theory?,” in “Handbook of Economic Growth,” Vol. 2 of *Handbook of Economic Growth*, Elsevier, 2014, chapter 1, pp. 515–563.
- Akcigit, Ufuk and Qingmin Liu, “The Role Of Information In Innovation And Competition,” *Journal of the European Economic Association*, August 2016, 14 (4), 828–870.
- and William R. Kerr, “Growth Through Heterogeneous Innovations,” *Journal of Political Economy*, forthcoming 2018.
- , Douglas Hanley, and Nicolas Serrano-Velarde, “Back to Basics: Basic Research Spillovers, Innovation Policy and Growth,” December 2016. University of Chicago manuscript.
- , Murat Alp Celik, and Jeremy Greenwood, “Buy, Keep, or Sell: Economic Growth and the Market for Ideas,” *Econometrica*, 2016, 84 (3), 943–984.

- Bartelsman, Eric J. and Wayne Gray, "The NBER Manufacturing Productivity Database," Working Paper 205, National Bureau of Economic Research October 1996.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," *American Economic Review*, February 2013, 103 (1), 305–34.
- Beintema, Nienke, Gert-Jan Stads, Keith Fuglie, and Paul Heisey, "ASTI Global Assessment of Agricultural R&D Spending," Technical Report, International Food Policy Research Institute 2012.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen, "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 07 2013, 81 (4), 1347–1393.
- Blundell, Richard, Rachel Griffith, and John Van Reenen, "Market Share, Market Value and Innovation in a Panel of British Manufacturing Firms," *The Review of Economic Studies*, 1999, 66 (3), 529–554.
- Branstetter, Lee G, "Are Knowledge Spillovers International or Intranational in Scope?: Microeconomic Evidence from the U.S. and Japan," *Journal of International Economics*, 2001, 53 (1), 53–79.
- Congressional Budget Office, "Research and Development in the Pharmaceutical Industry," Technical Report October 2006.
- Cowen, Tyler, *The Great Stagnation: How America Ate All The Low-Hanging Fruit of Modern History, Got Sick, and Will (Eventually) Feel Better: A Penguin eSpecial from Dutton*, Penguin Group US, 2011.
- Dinopoulos, Elias and Constantinos Syropoulos, "Rent Protection as a Barrier to Innovation and Growth," *Economic Theory*, 2007, 32 (2), 309–332.
- and Peter Thompson, "Schumpeterian Growth without Scale Effects," *Journal of Economic Growth*, December 1998, 3 (4), 313–335.
- Fernandez-Cornejo, Jorge et al., "The seed industry in US agriculture: An exploration of data and information on crop seed markets, regulation, industry structure, and research and development," Technical Report, United States Department of Agriculture, Economic Research Service 2004.
- Flamm, Kenneth, "Has Moore's Law Been Repealed? An Economist's Perspective," *Computers in Science and Engineering*, forthcoming.

- Food and Drug Administration, "Summary of NDA Approvals & Receipts, 1938 to the present," 2013. Online data report.
- Fuglie, Keith, Paul Heisey, John L King, Kelly Day-Rubenstein, David Schimmelpennig, Sun Ling Wang, Carl E Pray, and Rupa Karmarkar-Deshmukh, "Research investments and market structure in the food processing, agricultural input, and biofuel industries worldwide," *USDA-ERS Economic Research Report*, 2011, (130).
- Garcia-Macia, Daniel, Chang-Tai Hsieh, and Peter J. Klenow, "How Destructive is Innovation?," April 2016. Unpublished manuscript.
- Gordon, Robert J., *The Rise and Fall of American Growth: The US Standard of Living since the Civil War*, Princeton University Press, 2016.
- Griliches, Zvi, "Productivity, R&D and the Data Constraint," *American Economic Review*, March 1994, 84 (1), 1–23.
- Grossman, Gene M. and Elhanan Helpman, *Innovation and Growth in the Global Economy*, Cambridge, MA: MIT Press, 1991.
- Ha, Joonkyung and Peter Howitt, "Accounting for trends in productivity and R&D: A Schumpeterian critique of semi-endogenous growth theory," *Journal of Money, Credit and Banking*, 2007, 39 (4), 733–774.
- Hall, Bronwyn H. and Jacques Mairesse, "Exploring the relationship between R&D and productivity in French manufacturing firms," *Journal of Econometrics*, 1995, 65, 263–293.
- Hopenhayn, Hugo A., "Entry, Exit, and Firm Dynamics in Long-Run Equilibrium," *Econometrica: Journal of the Econometric Society*, 1992, pp. 1127–1150.
- Howitt, Peter, "Steady Endogenous Growth with Population and R&D Inputs Growing," *Journal of Political Economy*, August 1999, 107 (4), 715–730.
- Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *Quarterly Journal of Economics*, 2009, 124 (4), 1403–1448.
- Huffman, Wallace E and Robert E Evenson, *Science for agriculture: A long-term perspective*, John Wiley & Sons, 2006.
- Huffman, Wallace E., Yu Jin, and Zheng Xu, "The economic impacts of technology and climate change: new evidence from U.S. corn yields," January 2018. Working Paper.

- Jaffe, Adam B., "Technological Opportunity and Spillovers of R&D: Evidence from Firms' Patents, Profits and Market Value," *American Economic Review*, December 1986, 76 (5), 984–1001.
- Jones, Benjamin F., "The Burden of Knowledge and the Death of the Renaissance Man: Is Innovation Getting Harder?," *Review of Economic Studies*, 2009, 76 (1).
- Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy*, August 1995, 103 (4), 759–784.
- , "Growth and Ideas," in Philippe Aghion and Steven A. Durlauf, eds., *Handbook of Economic Growth*, New York: North Holland, 2005, pp. 1063–1111.
- and Jihee Kim, "A Schumpeterian Model of Top Income Inequality," *Journal of Political Economy*, forthcoming 2018.
- and John C. Williams, "Measuring the Social Return to R&D," *Quarterly Journal of Economics*, November 1998, 113, 1119–1135.
- Klette, Tor Jakob and Samuel Kortum, "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, October 2004, 112 (5), 986–1018.
- Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman, "Technological Innovation, Resource Allocation, and Growth," September 2015. Unpublished manuscript.
- Kortum, Samuel S., "Research, Patenting, and Technological Change," *Econometrica*, 1997, 65 (6), 1389–1419.
- Kremer, Michael, "Population Growth and Technological Change: One Million B.C. to 1990," *Quarterly Journal of Economics*, August 1993, 108 (4), 681–716.
- Kruse-Andersen, Peter K., "Testing R&D-Based Endogenous Growth Models," 2016. University of Copenhagen manuscript.
- Laincz, Christopher and Pietro Peretto, "Scale effects in endogenous growth theory: an error of aggregation not specification," *Journal of Economic Growth*, September 2006, 11 (3), 263–288.
- Li, Chol-Won, "Endogenous vs. Semi-endogenous Growth in a Two-R&D-Sector Model," *Economic Journal*, March 2000, 110 (462), C109–C122.
- Lichtenberg, Frank R., "The Impact of Public and Private Research Support on Premature Cancer Mortality and Hospitalization in the U.S., 1999-2013," Working Paper 23241, National Bureau of Economic Research March 2017.

- Lucas, Robert E., Jr., "On the Size Distribution of Business Firms," *Bell Journal of Economics*, 1978, 9, 508–523.
- Lucking, Brian, Nicholas Bloom, and John Van Reenen, "Have R&D Spillovers Changed?," 2017. Stanford University, unpublished manuscript.
- Melitz, Marc J., "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 2003, 71 (6), 1695–1725.
- Ngai, Rachel and Roberto Samaniego, "Accounting for Research and Productivity Growth Across Industries," *Review of Economic Dynamics*, July 2011, 14 (3), 475–495.
- Nordhaus, William D., "Do Real Output and Real Wage Measures Capture Reality? The History of Lighting Suggests Not," in Timothy F. Bresnahan and Robert J. Gordon, eds., *The Economics of New Goods*, University of Chicago Press, 1997, pp. 29–66.
- Oeppen, Jim and James W. Vaupel, "Broken Limits to Life Expectancy," *Science*, May 2002, 296 (5570), 1029–1031.
- Pardey, Philip G., Connie Chan-Kang, Jason M. Beddow, and Steven P. Dehmer, "Shifting Ground: Food and Agricultural R&D Spending Worldwide, 1960-2011," Technical Report, International Science and Technology Practice and Policy (InSTePP) Center, University of Minnesota 2016.
- Peretto, Pietro, "Technological Change and Population Growth," *Journal of Economic Growth*, December 1998, 3 (4), 283–311.
- , "A Note on the Second Linearity Critique," 2016. Duke University manuscript.
- , "Robust Endogenous Growth," 2016. Duke University manuscript.
- Perrin, Richard K, KA Kunnings et al., "Some effects of the US Plant Variety Protection Act of 1970.," *North Carolina State University. Dept. of Economics and Business. Economics research report (USA). no. 46.*, 1983.
- Pharmaceutical Research and Manufacturers of America, *2016 Biopharmaceutical Research Industry Profile*, Washington, DC: PhRMA, 2016.
- Pillai, Unni, "Vertical Specialization and Industry Growth: The Case of the Semiconductor Industry," 2016. work in progress.
- Romer, Paul M., "Endogenous Technological Change," *Journal of Political Economy*, October 1990, 98 (5), S71–S102.

- Segerstrom, Paul, “Endogenous Growth Without Scale Effects,” *American Economic Review*, December 1998, 88 (5), 1290–1310.
- Solow, Robert M., “Technical Change and the Aggregate Production Function,” *Review of Economics and Statistics*, August 1957, 39 (3), 312–320.
- Traxler, Greg, Albert KA Acquaye, Kenneth Frey, and Ann Marie Thro, “Public sector plant breeding resources in the US: Study results for the year 2001,” *USDA Cooperative State Research, Education and Extension Service*, 2005, pp. 1–7.
- University of York, “The Essential Chemical Industry,” 2016. [Online at <http://www.essentialchemicalindustry.org/>; accessed 8-September-2016].
- U.S. Department of Agriculture National Agricultural Statistics Service, 2016. [Online at <https://quickstats.nass.usda.gov/>; accessed 8-September-2016].
- U.S. Department of Agriculture National Institute of Food and Agriculture Current Research Information System, 2016. [Online at <http://cris.nifa.usda.gov/fsummaries.html>; accessed 8-September-2016].
- Vaupel, James W. and Vladimir Canudas-Romo, “Decomposing Change in Life Expectancy: A Bouquet of Formulas in Honor of Nathan Keyfitz’s 90th Birthday,” *Demography*, 2003, 40 (2), 201–216.
- Young, Alwyn, “Growth without Scale Effects,” *Journal of Political Economy*, February 1998, 106 (1), 41–63.
- Zambrowicz, Brian P. and Arthur T. Sands, “Knockouts Model the 100 Best-Selling Drugs — Will they model the next 100?,” *Nature Reviews Drug Discoveries*, January 2003, 2 (1), 38–51.
- Zeira, Joseph, “Innovations, Patent Races and Endogenous Growth,” *Journal of Economic Growth*, Jun 2011, 16 (2), 135–156.

A 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 3.4. 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 the main text in Table 6, but they are substantially negative nearly everywhere we look.

Table A1: Robustness Results for Research Productivity: $\lambda = 0.75$

Scope	Time Period	Average annual growth rate	Half-life (years)
Aggregate economy	1930–2015	-4.0%	17
Moore's law	1971–2014	-5.1%	14
Semiconductor TFP growth	1975–2011	-4.0%	18
Agriculture, US R&D	1970–2007	-3.2%	21
Agriculture, global R&D	1980–2010	-5.1%	14
Corn, version 1	1969–2009	-7.9%	9
Corn, version 2	1969–2009	-5.2%	13
Soybeans, version 1	1969–2009	-5.4%	13
Soybeans, version 2	1969–2009	-3.2%	22
Cotton, version 1	1969–2009	-1.9%	37
Cotton, version 2	1969–2009	+1.6%	-44
Wheat, version 1	1969–2009	-5.0%	14
Wheat, version 2	1969–2009	-2.9%	24
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
Compustat, market cap	3 decades	-7.8%	9
Compustat, employment	3 decades	-12.8%	5
Compustat, sales/emp	3 decades	-3.9%	18

Note: This table shows results paralleling those in Table 6 when we allow for 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 $\lambda = 1$. See the online data appendix for more details.