

NBER WORKING PAPER SERIES

A CALCULATION OF THE SOCIAL RETURNS TO INNOVATION

Benjamin F. Jones
Lawrence H. Summers

Working Paper 27863
<http://www.nber.org/papers/w27863>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2020

We thank Pierre Azoulay, Stefan Bechtold, Dietmar Harhoff, Chad Jones, Monika Schnitzer, Scott Stern, Manuel Trajtenberg, Martin Watzinger, seminar participants at ETH Zurich, the Max Planck Institute for Innovation and Competition, and conference participants at the NBER Role of Innovation and Entrepreneurship in Economic Growth conference for helpful comments. We gratefully acknowledge support from the Alfred P. Sloan Foundation under award G-2015-14014 and the AFOSR Minerva award FA9550-19-1-0354. All errors are our own. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2020 by Benjamin F. Jones and Lawrence H. Summers. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

A Calculation of the Social Returns to Innovation
Benjamin F. Jones and Lawrence H. Summers
NBER Working Paper No. 27863
September 2020
JEL No. O3,O4

ABSTRACT

This paper estimates the social returns to investments in innovation. The disparate spillovers associated with innovation, including imitation, business stealing, and intertemporal spillovers, have made calculations of the social returns difficult. Here we provide an economy-wide calculation that nets out the many spillover margins. We further assess the role of capital investment, diffusion delays, learning-by-doing, productivity mismeasurement, health outcomes, and international spillovers in assessing the average social returns. Overall, our estimates suggest that the social returns are very large. Even under conservative assumptions, innovation efforts produce social benefits that are many multiples of the investment costs.

Benjamin F. Jones
Northwestern University
Kellogg School of Management
Department of Management and Strategy
2211 Campus Drive
Evanston, IL 60208
and NBER
bjones@kellogg.northwestern.edu

Lawrence H. Summers
Harvard Kennedy School of Government
79 JFK Street
Cambridge, MA 02138
and NBER
lhs@harvard.edu

I. Introduction

Standards of living in advanced economies have risen dramatically over the last two centuries, with U.S. income per-capita currently 25 times its level in 1820 (Council of Economic Advisers 2011). Scientific and technological advances, ultimately delivering valuable new products and services, are thought to be critical drivers of these gains (Solow 1956, Mokyr 1990). Innovative advances also appear central to improving human health and life expectancy (e.g., Cutler et al. 2006). Yet measuring the social returns to scientific and technological advance has proven difficult. The challenge lies in the many spillover margins that appear inherent in the innovation process and the diffuse manner by which the fruits of research investments are often realized (e.g., Hall et al. 2010).

This paper does three things. First, it introduces a new method for calculating the average social returns to innovation. This method integrates across the many types of spillovers that innovative investments create. Second, this paper considers how the social returns vary according to potentially important but not commonly addressed features of innovation. These features include the roles of diffusion delays, capital embodiment, learning-by-doing, productivity mismeasurement, health outcomes, and international spillovers. The robust finding is that the social returns to innovative investments appear large. If a narrow set of innovative efforts (such as formal R&D) drive the bulk of productivity gains, then the social returns to these investments are enormous. If a much broader set of innovative efforts drives productivity gains, then the social returns to these broader activities appear merely large. In light of the high social returns, the final part of the paper discusses the prospects for increased innovative effort to accelerate improvements in standards of living and economic growth.

The existing literature emphasizes that the social gain from a new idea may differ substantially from the private gain captured by the original innovator (e.g., Griliches 1992, Nordhaus 2004). The divergence between private and social returns follows from various spillovers that appear integral to the innovation process. Positive innovation spillovers may include benefits to users (e.g., Trajtenberg 1989), benefits to imitators (e.g. Segerstrom 1991), and intertemporal benefits where new ideas enable additional innovations in the future (e.g., Romer 1990, Scotchmer 1991, Weitzman 1998). One can look to examples like electricity, the computer, or the Human Genome project -- and the new products, businesses, and industries they have spurred -- to

realize that the private gains to the initial innovators may be small compared to the productivity or health gains that result for society as a whole. However, while such spillovers suggest that the social returns of an innovation may substantially exceed the private returns, other forces may lead innovators toward over-investing in new ideas. Overinvestment could occur through business stealing (e.g., Aghion and Howitt 1998), research duplication (e.g., Dixit 1988), and/or intertemporal costs where finding new ideas today raises costs for finding new ideas later (e.g., Kortum 1997, Jones 2009).

In light of these spillovers, researchers have long been interested in understanding the social returns to innovation, with an emphasis on formal R&D investment. Case studies of specific technologies (e.g., Griliches 1958, Mansfield et al. 1977, Tewsbury et al. 1980) have counted up the R&D investments targeting specific products and then examined the benefits from the technologies developed. Other literature uses regressions to examine how firm and industry R&D investment pays off in productivity gains (e.g., Hall et al. 2010, Bloom et al. 2013). These regressions study spillovers by linking a given firm or industry's productivity to the R&D performed by other firms or industries. Regression methods have also been deployed at the national level to study how aggregate productivity gains are associated with aggregate R&D investment, including spillovers from R&D in other countries (e.g., Coe and Helpman 1995). Finally, macroeconomic growth models have been calibrated to data, and the social returns to R&D calibrated under various assumptions about functional forms and their parameter values (e.g., Jones and Williams 1998). These varied approaches typically reach broadly similar conclusions: the social returns to R&D are large.

At the same time, each of the above approaches faces methodological difficulties. Case studies of specific technologies raise the question of whether the results generalize to other technologies. This concern is more acute if the case studies tend to “pick winners”, which would lead to overstatement of the typical R&D returns. The regression methods often face challenges of causative interpretation. Further, regression methods must delineate the scope of spillovers, and distant spillovers or intertemporal spillovers are largely ignored. For example, these methods do not incorporate the role of basic research and the widespread but typically delayed influences that basic research may have in opening new avenues of commercial application (Ahmadpoor and Jones 2017). Yet it is exactly the innovations with diffuse implications – electricity, the

computer, genetics research, machine learning – that may be especially important for society and for understanding the returns to innovation investments.

Given these challenges, this paper introduces a new and complementary methodology. We present new calculations for the social returns to innovation investment, building on core features of the innovation and growth literatures. Our measures emphasize the advantages of examining the path of GDP, which acts to aggregate and net out complicated spillovers involved in the innovation process. The approach offers a seemingly quite general means of estimating the average social returns. Moreover, the simplicity of the method allows us to transparently examine the influence of other, potentially key features that are not typically addressed in studies of the social returns to R&D. These features include embodied vs. disembodied technological progress, diffusion rates, learning-by-doing, productivity mismeasurement, health benefits, cross-country spillovers, and other dimensions for assessing the social returns.

The intuition for our approach is straightforward. Modern growth theory, following the work of Robert Solow, tells us under reasonably broad conditions that the growth rate of GDP per person will be equivalent to the growth rate in total factor productivity (Solow 1956). In the absence of this productivity growth, per-capita income will remain constant. In advanced economies, long-run growth in productivity is often interpreted to come from investment in new ideas, which is the basis of modern endogenous growth theory (e.g., Romer 1990, Aghion and Howitt 1992).

Taking this approach seriously, the average returns to innovative investments are determined by linking the aggregate cost of innovation investments to the aggregate production increase that results. Intuitively, by looking at the net value-added gains in the GDP path, one can implicitly net out the spillover margins. By looking at total innovation investment, one includes both research successes and failures. A simple social returns calculation can proceed as follows. Let income per capita be y , innovation investment per capita be x , and the discount rate be r . If a year's worth of innovation investments creates a g percent increase in productivity, then the ratio of benefits to costs is:

$$\rho = \frac{g/r}{x/y}$$

The key idea here, as in endogenous growth theory, is that, by investing a GDP share x/y in innovation today (i.e., once), we permanently raise productivity in the economy by g percent, the present value of which is g/r . Notably, this approach suggests that the average social returns to innovation may be enormous. For example, if we take an R&D investment orientation, with the R&D share of GDP at its usual level in the U.S., $x/y \approx 2.7\%$, and let these investments drive productivity growth, then we have $g \approx 1.8\%$.¹ Standard discount rates then imply that 1\$ of R&D investment today on average creates over \$10 of economy-wide benefits in today's dollars.² This return is extremely large, but it follows from the basic mechanics of growth, as understood in advanced economies. That is, a permanent gain in living standards from a seemingly small investment in innovation will, by the above logic, tend to suggest enormous returns.

Having established this baseline, "R&D only" analysis, the paper examines several reasons it may be too high. First, we consider the role of diffusion, where the gains from R&D may pay off slowly, delaying the benefits that are achieved and thus reducing their present value. Second, we consider the role of capital deepening in accounting for some of the productivity gains and, related, consider the role of capital-embodied technical change, where the value of R&D investments may only be realized through investments in new types of fixed assets. Third, we consider the possibility that productivity growth occurs without formal R&D but due to other kinds of activities, such as new-venture creation or learning-by-doing. In each case, we calibrate adjusted returns. All of these analyses act to reduce the estimated social returns to innovation investments, but we will also argue that under plausible assumptions the returns still appear very high.

The paper then examines several reasons that, under the above calculation, the (already high) social returns estimates may be too low. We first consider the role of inflation bias, causing real GDP growth to understate gains due to product improvements and new product introductions. Second, we consider health, which is a main target of R&D investments and may bring large

¹ We will consider the role of capital deepening and many other extensions to this simple calculation in the body of the paper.

² An alternative social returns calculation is the internal rate of return. This is the discount rate at which the benefits and costs are the same. This rate of return is $r^* = \frac{g}{x/y}$ for the simple calculation above. Using the same values for g and x/y , we then have $r^* = 67\%$.

social returns, but where mortality and morbidity are not well accounted for in standard GDP per capita measures. Third, we consider international diffusion, where economies around the world may also benefit from the innovation investments undertaken in frontier economies.

Finally, the paper considers the distinction between the average and the marginal social return to innovation investments. Our calculations throughout are explicitly about the average return, which avoids having to assume very specific production functions. However, policymakers may naturally be more interested in the marginal investment returns. That is, policy choices will hinge on whether additionally increasing innovation investment levels will see the same kinds of returns that the average return calculations indicate. We therefore consider how to bridge between marginal and average return and present specific estimates of the marginal returns.

The rest of the paper proceeds as follows. In Section 2, we introduce our methodology and consider baseline calculations of the average social return. In Section 3, we consider reasons that these baseline calculations may be too high, and then consider reasons the baseline calculation may be too low. In Section 4, we consider distinctions between the average and the marginal social return, and discuss the prospects for increased innovation investment to raise the rate of advance in socioeconomic prosperity. Section 5 concludes.³

II. The Average Social Returns to R&D: A Baseline

In this section we introduce a baseline calculation of the average social returns to innovation investment. This method is meant to achieve three things. First, it acts to integrate across the many spillovers inherent to innovation. Second, it clarifies the basic logic for why the social returns to innovation appear high. Third, it provides a foundation for discussing, and clarifying, a range of additional and potentially first-order issues that bear on the social returns, which we will consider in Section 3.

A. Towards the Social Returns to Innovation

³ Three appendices provide further context and results. Appendix A details the many types of spillovers that appear inherent to the innovation process and provides examples of each. Appendix B reviews the existing empirical literature that works to confront these spillovers and estimate the social rate of return. Appendix C provides proofs of the formal results in the paper.

The social returns to innovation depend on the cost of innovation and the benefit that results. That the social returns tend to appear high in various analyses⁴ speaks to the fact that the costs often appear low yet the gains often appear substantial and durable. In particular, a new idea, method, design, etc., can be created at some upfront cost (i.e., paid once) but then raise productivity more-or-less forever. For example, calculus, invented in the 17th century, was a permanent advance in mathematics that has been used ever since.

This idea – that an upfront cost leads to permanent advance – applies conceptually to innovation as a whole. For example, consider a software innovation that raises the productivity of workers by p_1 percent. If this were the last innovation ever produced, it would continue to provide this p_1 percent gain forever. But let's say instead that another software innovation comes along that replaces the original software and raises the productivity of workers by an additional p_2 percent. We can think of this sequence of innovations in two ways. First, we can think of the original innovation as producing a permanent gain of p_1 percent and the second innovation as, at some further innovation cost, producing an additional gain of p_2 percent. In this sense, the marginal gain from the original innovation remains. Alternatively, we might consider the average return to both together, adding up the innovation costs and adding up the total productivity gains so that the innovation investments have a permanent effect in combination.

This kind of thought experiment lies behind the “case study” approach. Because innovations interact in complex ways, and many small innovations may together advance productivity in a given product line, separating out the marginal returns of each innovation is difficult. The case study method thus often pools the innovation costs and benefits across many related innovations and calculates an average social return to the broader technological advance, rather than the marginal return of each micro-innovation (e.g., Griliches 1958).

The limitation of case studies is one of representativeness, where they are unlikely to describe innovation investment returns in general. In particular, case studies of failures are rare, even though failure in innovation is common (Arrow 1960; Kerr et al. 2014). By leaving out failures, case studies may overstate the general social returns to R&D. Yet the advantages of the case study approach may still be had, separately from this limitation, by expanding the boundaries of

⁴ See Appendix B for a review.

the exercise. This occurs if one applies the approach to the economy as a whole. By aggregating across all innovation investment, one incorporates not only successful investments but also the “dry holes” of failed investments. Total innovation costs also incorporate the potential wasteful duplication of innovative efforts. On the benefit side, the path of aggregate productivity gains nets out the imitative and business stealing spillovers between and across firms and industries. The long-run path of productivity gains also accounts for intertemporal spillovers and the benefits of basic research.

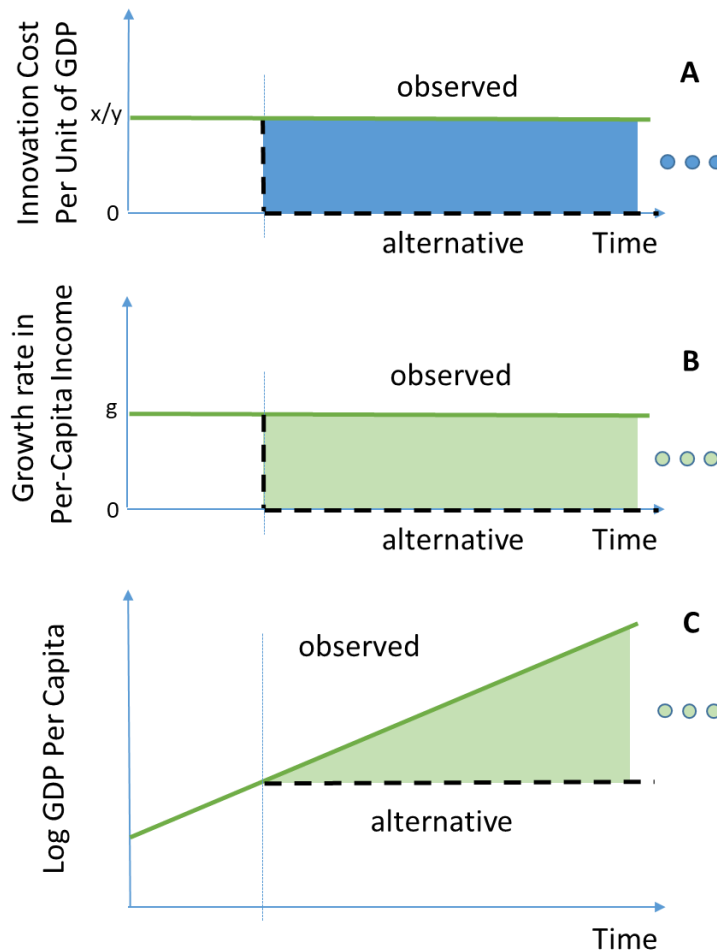


Figure 1: Conceptual Model of the Economy-Wide Returns to Innovation Investment

B. The Average Social Returns to Innovation

To develop this idea formally, consider two thought experiments. Both provide a baseline view of the social returns to innovation (which we extend further below). In the first thought

experiment, we “turn off” all innovative investments forever and consider the costs and benefits. See Figure 1. On the innovative investment side, we move from the observed level of innovation investments to no innovation investment (Figure 1A). On the output side, if productivity advances cease with no further innovation, then we move from the observed level of growth to a state of no further growth (Figure 1B). Thus per-capita income remains constant (Figure 1C). This alternative, no growth state is also the outcome in modern endogenous growth theory (e.g., Romer 1990, Aghion and Howitt 1992), but note that the key assumption here is more general than specific endogenous growth models.⁵

The average social returns to the observed level of innovative investments then follows directly. The cost of innovation is the present discounted value of the innovation investments. This is the present value of the rectangle in Figure 1A. The benefits of innovation are the present value of the discounted gain in per-capita income. This is the present value of the triangle in Figure 1C. The ratio of the benefits to the costs then gives the social cost-benefit ratio.

As shown in the appendix, the social-cost benefit ratio is extremely simple and intuitive. It is

$$\rho = \frac{g/r}{x/y} \quad (1)$$

The cost, in the denominator, is the ratio of innovation investment expenditure (x) to GDP (y). The benefit, in the numerator, is the growth rate (g) that results, discounted to the present at the discount rate (r). We are suppressing time in this expression to emphasize that the ratio of innovation investment expenditure to GDP and the growth rate of income are approximately constant over time.

This expression, although derived over the entire time path of innovation expenditure and the entire time path of productivity gains, produces an interpretation based on the intuitive nature of innovation gains. Namely, we can think of the cost benefit ratio as the cost of one year’s innovation (x/y) producing a stream of output gains that are g percent higher. The present value of this permanent output gain is g/r .

⁵ That is, if there are diminishing returns to capital, then the absence of productivity growth means no growth in per-capita income.

As an alternative thought experiment, consider Figure 2. Here we imagine that we turn off innovation investment for one year only (Figure 2A) rather than forever (Figure 1A). Since we do not innovate that year, we see no gain in productivity growth that year (Figure 2B and 2C). However, at the end of the year, we start innovating again. In particular, we undertake exactly the same innovation projects as on the observed path. It should therefore be clear that the economy, with exactly the same innovations, leads to exactly the same productivity levels, only now the innovation costs occur with a one year delay and the economy arrives at each productivity level one year later.




Figure 2: Alternative Conceptual Model of the Economy-Wide Returns to Innovation Investment

In this alternative thought experiment, the present value of the innovation costs on the observed path versus the alternative path is one year's innovation costs, or x/y . The benefits are being g

percent richer in each future period, the present value of which is g/r . These are the present values of the shaded areas in Figure 2A (costs) and Figure 2C (benefits). The social benefit-cost ratio is then, again, exactly as in (1).

Finally, an alternative calculation is to measure the discount rate where the costs and benefits would be equal (i.e., where $\rho = 1$). This internal rate of return is written

$$r^* = \frac{gy}{x} \quad (2)$$

providing a social internal rate of return, as an alternative to the social benefit-cost ratio.

C. The Average Social Returns: An “R&D Only” Baseline

With the expression (1), we can now calculate a baseline average social return to innovative investments. Taking the U.S. economy, we have $g = 1.8\%$ as the average long-run growth rate. Using total public and private R&D investment to account for innovation investments (x), the long-run average value of x/y is approximately 2.7%.⁶ Taking a discount rate of 5%, we then have

$$\rho = \frac{.018/.05}{.027} = 13.3$$

This says that \$1 of R&D investment today produces, on average, a \$13.3 benefit in today’s dollars. In other words, this baseline calculation suggest that the social returns to R&D are enormous.

An open question here is what discount rate to choose. The lower the discount rate, the greater the innovation benefits. The social discount rates used by governments range from 3.5% (U.K.) to 7% (U.S.). Some argue that social discount rates should be lower still, and equilibrium real interest rates have been trending downwards for decades (Rachel and Summers 2019). U.S. 30-year inflation-protected government bonds point to a 1% discount rate on average over the last decade. Using such low discount rates would further amplify the social returns, but even high discount rates suggest the social returns are very large. See Table 1.

⁶ U.S. R&D expenditure is based on numerous surveys of the National Science Foundation, and include R&D performed and funded by private businesses with at least five employees, federal and state governments, universities, and non-profit organizations.

As an alternative calculation, focusing on the internal rate of return, (2), and assuming again that $g = .018$ and $x/y = .027$, we have $r^* = 67\%$. By this standard, the social returns are also enormous. For example, if a private citizen could access an investment with a 67% annual rate of return, that individual could become very rich very quickly. Of course, as a social return, this rate of return is not available to an individual investor. But it may be available to society as a whole. The question becomes whether and how society can further invest to take advantage of this high return. See Section 5.

Overall, this simple baseline calculation based on R&D expenditure suggests that the average social returns to innovation are very large. This result aggregates across failed and successful R&D projects. It also incorporates the manifold spillovers involved in R&D, including intertemporal spillovers. In magnitude, the result reinforces the common finding in the prior literature, surveyed in Section 2, and comes in at the upper range of existing estimates. We can now adjust this baseline calculation on numerous dimensions to assess whether it may be too large or too small, which is the subject of the balance of the paper.

Table 1: The Average Social Returns, by Social Discount Rate

Social discount rate (r)	Average Social Benefit-Cost Ratio (ρ)
1%	66.7
2%	33.3
3.5%	19.0
5%	13.3
7%	9.5
10%	6.7
67%	1

III. Extending the Baseline

As an overarching framework for what follows, consider the following adjustment to the social return calculation. We write

$$\rho = \beta \frac{g/r}{x/y} \tag{3}$$

where the new term, β , provides for an upward or downward adjustment to the social return. In this section, we first consider forces that make β less than one, so that the baseline calculation in Table 1 is too high. We will then consider forces that make β greater than one, so that the baseline calculation is too low.

A. Reasons the baseline social returns may be too high

1. Lags

The above baseline assumes that the payoff from R&D investments occurs immediately. Yet there may be substantive delays in receiving the fruits of R&D investments. Other things equal, the longer the delays until the benefits are realized, the lower the return.

A simple approach to potential delays assumes that R&D investments borne today increase productivity permanently starting D years in the future. This leads to a straightforward correction to the present value of the benefit stream. The calculation is the same as before, but we now must include a discount factor where⁷

$$\beta = e^{-\hat{r}D} \quad (4)$$

To make an explicit adjustment, we can consider various micro evidence. For businesses, the literature suggests a relatively short delay between R&D investment and product introduction. Mansfield (1971) finds a 3-year median delay. Ravenscraft and Scherer (1982) find, in a survey, that 45% of firms report a 1-2 year delay, 40% of firms report a 2-5 years, and only 5% of firms report a delay beyond 5 years. Pakes and Schankerman (1984) impute a delay of 1.5-2.5 years between R&D investment and first revenues. For Compustat firms, which focus on mature firms in mature industries, Argente et al. (2018) estimate a one-year delay between R&D and product introduction.

The first introduction of the product is not, however, the time at which use in the market peaks. Leonard (1971) studies 19 manufacturing industries and finds that the growth continues from the second year up until the ninth year after the R&D investment. In mature consumer sectors, the delays to market peak may be shorter. For example, Argente et al. (2019) find that the new consumer products from publicly-traded firms typically hit their sales peak one year after

⁷ The appropriate discount factor is $\hat{r} = r - g$. This accounts for both the discount rate, r , for future income as well as income growth, g , which expands the income over which today's innovation will ultimately be felt.

introduction. All told, studies of R&D, product introductions, and product sales suggest quite rapid linkages between up-front costs and peak market payoffs. A total delay of 3-6 years appears reasonable, and a 10-year delay appears very conservative.

For basic research, the delays are naturally longer. Using regression analysis, Adams (1990) suggests a 20-year lag between academic research and productivity growth in the relevant industry. One can also link specific patents to the underlying scientific research that each patent cites (Ahmadpoor and Jones 2017). Examining all U.S. patents indicates an average delay of 6 years from the patent application to its direct precursor science publications. To the extent that basic research pays off indirectly (i.e., basic research leads to further research that, eventually, becomes an input to marketplace inventions) citation network analysis suggests that even remote basic research investments begin paying off within 20 years.

Table 2 reconsiders the baseline social returns calculations using a range of delays.⁸ Aggregating across the different types of research, a middle-of-the-road delay estimate may be 6.5 years, and a conservative estimate would be 10 years.⁹ An extremely conservative estimate would be 20 years. Using any of these delays, the average social returns to R&D still appear very large. See Table 2. Even taking a very conservative 20-year average delay, which is well beyond what the micro-evidence indicates, one would still see a \$4.9 present value benefit for each \$1 spent on R&D.¹⁰ The internal social rate of return declines relatively sharply from the baseline value of $r^* = 67\%$ with extended delays, because high internal rates of return heavily discount the future. With a delay of 20 years, the internal social rate of return falls to 11%.

⁸ We use $\hat{r} = 5\%$ in these calculations of the social benefit-cost ratio. With a growth rate of $g = 1.8\%$, this value for \hat{r} assumes a high discount rate of $r = 6.8\%$. To the extent that appropriate social discount rates are lower than this discount rate, the social benefit-cost ratios in Table 2 are conservative.

⁹ The National Science Foundation reports (NSF 2019) that 62.5% of recent U.S. R&D spending represents product development (i.e., R&D targeted toward the development or improvement of specific products or processes), 20% of spending represents applied research (i.e., research that has a specific practical aim or objective), and the remaining 17% of spending represents basic research (i.e., without any particular application in view). Taking a mainstream estimate of a 3-year delay for product development R&D, a 6-year delay for applied R&D, and 20-year delay for basic R&D, the average delay (weighted across expenditure on each category) would be 6.5 years. Taking a conservative estimate of a 5-year delay for development R&D, a 10-year delay for applied R&D, and a 30-year delay for basic R&D, the average delay would be 10 years.

¹⁰ A more sophisticated version of this delay adjustment considers not a single delay, D , but rather uses the full distribution of delays in the micro literature. In practice, however, the more sophisticated approach leads to similar conclusions.

Table 2: The Average Social Returns, for Different Benefit Lags

Delay in Years (D)	Corrective Factor (β)	Average Social Benefit-Cost Ratio (ρ)	Average Social Rate of Return (r^*)
0	1	13.3	67%
3	0.86	11.5	29%
5	0.78	10.4	23%
6.5	0.72	9.6	20%
10	0.61	8.1	16%
20	0.37	4.9	11%

2. Incorporating Capital Investment

The baseline approach assumes that growth relies on the innovative investment, x . This approach follows from the standard neoclassical idea that total-factor-productivity gains are necessary to achieve positive steady-state growth. Hence, one might accrue the benefit of growth to R&D. However, this approach implicitly ignores potentially important features, including the contribution of capital investment. Here we will introduce capital investment explicitly into the observed and counterfactual growth scenarios.

One can incorporate the capital investment part of productivity growth – or “capital deepening” – under two different viewpoints. In one viewpoint, technological progress is “disembodied” from capital inputs. In the second viewpoint, new technologies must be embodied in new capital inputs, which bring additional costs. Consider each perspective in turn.

Disembodied Productivity Growth

In the disembodied perspective, the productivity gains from innovation are felt independently of capital investment. We can then parse per-capita income gains into two features. First is the direct (disembodied) gain from innovation that occurs holding capital fixed; the second would be the capital deepening that these productivity gains further spark.

In standard neoclassical growth theory, where production is Cobb-Douglas, growth in per-capita income follows as

$$g_y = \alpha g_k + (1 - \alpha)g_A \quad (5)$$

where g_y is the growth in per-capita income, g_k is the growth in capital per person (capital deepening), and g_A represents technological progress. The term α is the capital share of income, which empirically is approximately one-third. Thus, if technological progress follows from innovative investments and these gains are felt in a fashion disembodied from capital investments, then we have a very simple correction to our baseline. Namely,

$$\beta = 1 - \alpha \approx \frac{2}{3}$$

It is then a simple matter to revise the estimates in Table 1. Simply multiply the social benefit-cost ratio by two-thirds. The returns to innovative investments with this correction still appear very large. The same is true if one applies this correction to Table 2, which further accounts for potential delays.

Capital-Embodied Productivity Growth

Alternatively, we may believe that the economic gains from innovative investments are in large part realized through the embodiment of these ideas in new forms of capital. For example, an innovation in microprocessors is useful only if it is built into microprocessors themselves. The same may be true for myriad forms of innovation that are embodied in capital equipment and structures. If so, then the benefits from new ideas (and advances in standards of living) require both the R&D expenditure and the investment in building the new or improved capital inputs.

In short, the costs are more than just the R&D. We can no longer make a clean separation between the innovation and capital deepening components, and the natural correction here is to include both investments together. The added piece is the cost of capital deepening. One can proceed empirically here and also consider theoretical bounds on these additional costs.

Empirically, since 1960, the annual net domestic investment of the U.S. private sector has averaged 4.0% of GDP. This net investment, which does not include R&D expenditure, incorporates capital deepening. Viewing these costs as necessary to realize the gains of R&D, the total costs to innovative investments would then be viewed as the summation of the R&D investment (2.7%) and net domestic investment (4.0%), or 6.7% of GDP. Thus we could say empirically that

$$\beta \approx \frac{2.7}{6.7} = 0.40$$

This “embodied” version of capital deepening thus reduces the social returns to R&D more than a “disembodied” innovation calculation. Nonetheless, given the baseline social returns in Table 1, such a correction still points to extremely high social returns. See Table 3.¹¹

The above calculation can be further adjusted in several senses. First, net investment costs are not just about capital deepening, but also include investment costs that extend the capital stock over a growing population. The U.S. population has grown annually at 1.0% on average since 1960, and the U.S. workforce has grown at 1.5% annually. This makes the above adjustment conservative – tending to understate the social returns. Second, capital deepening might occur outside the net investment data from domestic businesses, including, e.g., through infrastructure or other public investments. While private capital equipment investment may be especially important, other investment costs, if unaccounted for, would make the above adjustment less conservative.

We can generalize as follows. Along the equilibrium growth path of the economy, the deepening component of investment is equivalent to the growth rate in per-capita income times the capital-output ratio. That is,

$$\frac{i_{deep}}{y} = \frac{k}{y} g \quad (6)$$

where i_{deep} is the cost of investment that increases capital per worker and the other terms are defined above. The ratio of the U.S. capital stock to U.S. GDP has averaged 3.5 since 1960. Thus, capital deepening costs are approximately 3.5g. This suggests that the capital deepening cost would, in total, be 3.5 x 1.8%, or 6.3% of GDP. This correction is larger than the 4.0% net domestic investment cost taken from private businesses. Adding in the 2.7% of GDP for R&D, the total cost of innovation (idea creation and implementation) rises to 9.0% of GDP. This suggests that $\beta \approx 0.3$.

¹¹ For the capital deepening component, the issue of lags is substantially less germane. In the embodied-innovation perspective, the costs of the capital investment occur very close to the time of use of these embodied ideas in the economy. Thus the lag corrections in Table 2 do not apply to most of the innovation cost under this embodied-innovation perspective.

Table 3 summarizes these results. The conclusion is that accounting for capital deepening will reduce the social returns to innovation, and that embodied technical change reduces these returns more than disembodied technical change. Nonetheless, the main conclusion is that the social returns to innovation still appear extremely large.

Finally, note that incorporating capital investment doesn't diminish the society wide gains. Rather it acts to spread the gains over a broader set of investments, beyond R&D. The social returns to capital deepening thus appear much larger than the equilibrium private rate of return to capital investment would suggest. To the extent that embodiment is important, R&D investment and capital investment collectively unlock large social returns. From a policy point of view, supporting R&D and capital deepening together would may then be important to attain the high social returns from innovative investments.

Table 3: The Average Social Returns with Capital Costs

Capital Costs to Realize Productivity Gains	Corrective Factor (β)	Average Social Benefit-Cost Ratio (ρ)	Average Social Rate of Return (r^*)
	<i>Disembodied Technical Change</i>		
--	0.66	8.9	44%
	<i>Capital-Embodied Technical Change</i>		
Net Domestic Private Investment	0.40	5.3	27%
All Capital Deepening	0.30	4.0	20%

3. Other Innovation Costs

The above analysis links productivity growth in the economy to R&D investment and capital investment. To the extent that innovations come from other types of investment, one would undercount true innovation costs and thus overstate the social returns. Here we consider these possibilities.

One potentially important source of innovation is the effort of entrepreneurs. In particular, R&D expenditure measures in the U.S. explicitly do not include the innovative activities of businesses

with less than 5 employees.¹² Most of these businesses are not growth-oriented, representing self-employment or permanently small businesses like single-establishment restaurants, nail salons, etc.; however, a small set of new businesses are focused on creating transformative innovations (e.g., Guzman and Stern 2017, Azoulay et al. 2020). A practical adjustment to account for these innovative startups can include total venture capital investment as an additional innovation cost. Since 1995, total venture capital investment in the U.S. has been as high as \$130B (2018) but has often been less than \$30B per year.¹³ On average, total venture capital investment since 1995 is less than 0.3% of GDP. Adding all venture capital investment would therefore raise the costs of innovation investments from 2.7% to 3.0% of GDP, suggesting $\beta = 0.9$. This adjustment only modestly affects the social returns.

Additional inroads to estimating other innovation costs come from business surveys. Eurostat's Community Innovation Survey asks firms to compare their R&D costs with any additional innovation costs. In the 2016 survey, these numbers are available for 28 countries (Eurostat 2019). Taking these countries as a whole, firms report that R&D expenditures amount to 55% of total innovation costs. The non-R&D innovation costs are primarily investment in capital assets, including equipment, machinery, and software. These costs are linked to the adoption and diffusion of the innovations (Brouwer and Kleinknecht 1997, Evangelista et al. 1998) and can thus be seen as a component of the capital-embodied productivity gains analysis in Section IV.B above. Thus one could take $\beta = 0.55$ to account for these "other costs", or one can use the broader β correction that already encapsulates broad forms of associated capital investment costs, as in Table 3.

4. Learning-by-Doing and Incidental Innovation

A different challenge to the above estimates comes on the benefit side. To the extent that productivity growth comes from other sources, assigning the productivity growth to explicit R&D investment, new venture investment, and capital investment would overstate the social returns to these investments. What happens if new ideas, or the spark of ideas, come outside the above processes? Some innovative ideas may emerge from incidental inspirations among

¹² See National Science Foundation, Science and Engineering Indicators 2018, <https://www.nsf.gov/statistics/2018/nsb20181/>.

¹³ Source: Pitchbook/NVCA.

workers in the course of ordinary labor activities, rather than through focused investment expenditure. Learning-by-doing is typically seen as productivity improvements that come through accumulated experience and skill in a production process (Arrow 1961, Bessen 2015). Such advances might be seen as essentially “free” sources of productivity gains.

To relate this possibility to the social returns to innovation, consider three perspectives. First, canonical examples of learning-by-doing, like airframe manufacturing (Wright 1936) suggest that these productivity gains can be large. However, these gains typically hinge on and occur after the introduction of a new good or production process. In this sense, learning-by-doing acts as a kind of free innovation process that comes *ex-post* of necessary up-front costs, such as the R&D investment in the airframe design or the capital investment in producing the manufacturing machinery and facilities. From this perspective, learning-by-doing is then akin to the lag adjustment in Section 3.A.1, and one can therefore incorporate learning-by-doing by allowing for lags in the benefits of the up-front innovation investment costs.

Second, there may be “free ideas” that come to individuals, including individuals who do not participate in any measured R&D or investment process. Concretely, an individual may have an idea while driving to work, or while engaged in some work process, concerning things like a new jet engine design, computer application, web service, or medical device. Product users may also be an important source of new concepts (Von Hippel 1976). However, to the extent that the initial idea may appear to come for free (i.e., outside a measured investment process) the implementation of a new computer application, medical device, etc., will presumably bear further development costs and/or investments in fixed assets. If so, then using a broad measure of investment as featured in Section IV.B above, should still capture the total costs for achieving the productivity gain.

Finally, there may be free ideas that require no investment to implement. Returning to the learning-by-doing literature, the so-called “Horndal Effect” provides an example of productivity gains in a Swedish ironworks that appears to occur without any formal investment (Lundberg 1961). If such gains are actually investment-free, and responsible for a substantial portion of productivity growth, then the average returns to measured innovation investments above would be correspondingly lower. It is difficult to assess this possibility in general. The reader can adjust the social returns, however, in a straightforward manner, by choosing β as the share of

productivity growth that hinges on actual investment. Taking the broad returns in Table 3, one could assume that half of productivity gains are achieved without relying on any R&D investment, any new venture investment, and any capital investment, and still the average social returns to these measured innovation investments would be large.

5. Summary

Collectively, we have considered several independent reasons that the baseline social returns calculation may be too high. Analyzing each correction in isolation, the social returns to innovative investments tend to remain high. Analyzing several potential corrections at once, it is still difficult to find a result where the social returns are not high. The most important correction appears to be how we treat capital investment, especially if the results of R&D must be embodied in capital equipment. Incorporating long delays between the up-front investments and the ensuing productivity gains, which can additionally incorporate learning by doing, the social benefits still substantially exceed the costs. As we will discuss next, there are also several forces pushing in the other direction, which suggest that the baseline calculation may be too low.

B. Reasons the baseline calculation may be too low

1. Productivity growth mismeasurement

Calculating the path of real GDP is challenging. Economists have long recognized various problems in inflation statistics -- including substitution bias, product improvement, and the introduction of new goods -- which undermine the accuracy of inflation indices. Since at least the Stigler commission (Stigler et al. 1961), the economics consensus has been that inflation in consumer and producer prices is overstated, and hence real GDP growth is understated. The Boskin Commission found that the consumer price index overstated inflation rates by 1.10% per year, with a “plausible range” of 0.80% to 1.60% per year (Boskin et al. 1996). The most substantial source of bias (0.60%) found was due to the introduction of new goods and quality changes in existing goods – i.e., outcomes of innovation itself, so that the benefits of innovation are understated in concrete ways. The Boskin Commission’s findings and recommendations led to changes in price measurement approaches, and net of these changes the inflation bias was subsequently estimated to be an approximate 0.65% overstatement per year (Gordon 2000), although quality advances and new goods problems remain particularly challenging. The ongoing advance of computing, the Internet, and associated digital services has now led many

economists to believe that inflation bias may be much worse again today (Goolsbee and Klenow 2018, Brynjolfsson et al. 2019).

Applying these kinds of biases to gross domestic product overall will substantially increase the growth rate of the economy. The baseline social return to innovation will consequently increase. The correction is

$$\beta = 1 + \frac{\textit{inflation bias}}{g} \tag{7}$$

Table 4 considers corrections to the baseline social returns, under various assumptions about inflation bias. Taking the Boskin Commission’s central estimate (1.10%), which may be the right number historically, we see that $\beta = 1.6$. Even with the more modest Gordon estimate of inflation bias (0.65%), the social returns are elevated by more than one-third.

Table 4: The Average Social Returns, Correcting for Inflation Bias

Inflation Bias (% per annum)	Corrective Factor (β)	Average Social Benefit-Cost Ratio (ρ)	Average Social Rate of Return (r^*)
0.00	1	13.3	67%
0.40	1.22	16.3	81%
0.65	1.36	18.1	91%
0.80	1.44	19.3	96%
1.10	1.61	21.5	107%
1.60	1.89	25.2	126%

2. Health outcomes

A large portion of R&D, and related capital investment, is related to improving health and longevity. While health may influence productivity, a common target of health innovation is to extend life. Average life expectancy in the U.S. has risen considerably, from 47.3 in 1900 to 69.7 in 1960 to 78.5 in 2018, and there has been vast progress against infant mortality.¹⁴ Among the contributing factors to these health improvements, innovations play an important role,

¹⁴ Infant mortality in the U.S. has fallen from 100 deaths per 1,000 live births (1915) to 26 deaths per 1,000 live births (1960) to 5.6 deaths per 1,000 live births in 2017. As Murphy and Topel point out, in 1900 in the U.S. 18% of males did not reach their first birthday, but by 2005 one doesn’t reach an 18% mortality rate until age 62 (Murphy and Topel 2006).

including the advent and advance of vaccines, antibiotics, cardiovascular treatments, diagnostic and imaging technologies, surgical methods, and oncology products.

From a social return point of view, one could attempt to incorporate some portion of the health gains in the total benefits of innovation. Alternatively, one could remove health R&D from the cost side to produce an average return to non-health-related innovation. On the cost side, about 20-25% of government-financed R&D flows through the National Institutes of Health. In the private sector, approximately 18% of private-sponsored R&D in the U.S. in 2016 came from pharmaceuticals and medicines alone.¹⁵ A broader collective estimate calculates that total medical and health related R&D in the U.S. has risen from \$143 billion to \$182 billion from 2013 to 2017, amounting to approximately 30% of total U.S. R&D expenditure.¹⁶ Adjusting R&D to remove health expenditure will cause the baseline social return to (non-health-related) innovation to consequently increase. The correction is

$$\beta = \frac{x}{x - \text{health R\&D}} \quad (8)$$

Table 5 considers corrections for alternative measures of health R&D investment. Taking an estimate where health expenditures are 20% of total R&D, the social returns to non-health related innovation rise by 25%. This estimate is somewhat conservative in terms of potentially understating health R&D costs. But it may be non-conservative in that it assumes that these health benefits stand fully outside productivity increases; e.g., they support longevity or the health of the retired individuals, as opposed to being investments in worker productivity.

Adjusting instead on the benefit side, the most direct correction is to incorporate the value of living longer. This is a difficult calculation to make, even in principle. Economists often rely on the value of a statistical life, which can be based on observed expenditure to reduce the risk of death. Using this method and valuing life at \$3 million in 1990 dollars, Nordhaus (2005) finds that rising life expectancy has produced annual gains for individuals that exceeds the measured growth in the consumption of goods and services.¹⁷ Murphy and Topel (2006) further find

¹⁵ National Science Foundation 2018, Table 2, <https://nsf.gov/statistics/2018/nsf18312/>.

¹⁶ Source: Research America 2018 (https://www.researchamerica.org/sites/default/files/Policy_Advocacy/2013-2017InvestmentReportFall2018.pdf) and NSF Science and Engineering Indicators 2018 (<https://www.nsf.gov/statistics/2018/nsb20181/>).

¹⁷ Nordhaus uses a somewhat conservative value of life measure, compared to other studies and to U.S. government practice in performing cost-benefit analyses.

enormous welfare gains in the United States from both increased longevity and higher quality of health while alive. In their estimation, the social benefits from improved health are several times the increased expenditure on health care overall.

Table 5: The Average Social Returns, Accounting for Health

	Corrective Factor (β)	Average Social Benefit-Cost Ratio (ρ)	Average Social Rate of Return (r^*)
<i>Social returns for non-health-related innovation</i>			
Health R&D/Total R&D (%)			
0	1	13.3	67%
20	1.25	16.7	83%
25	1.33	17.8	89%
30	1.43	19.0	95%
<i>Social returns including benefits of life extension</i>			
Health Consumption Share (%)	Health Consumption Growth rate (%)		
25	1.0	1.19	15.8
30	1.5	1.36	18.1
40	2.0	1.74	23.2
50	2.5	2.39	31.9

Explicitly adjusting the social return calculation is somewhat difficult, however, because one needs not just the change in value from improved health but also a different baseline definition of real consumption that incorporates the value of being alive. That is, real GDP per capita is now higher in every period because of this “health consumption”, which we denote h . As we show in the appendix, an appropriate correction to the social return is

$$\beta = \left(1 + \frac{s_h}{1 - s_h} \frac{g_h}{g}\right) \quad (9)$$

where g_h is the growth rate in health consumption and s_h is the share of health consumption in the augmented GDP measure. This correction shows that adding positive growth in health consumption can only increase the social returns to innovation. The amplification factor is increasing in both the growth rate of this health consumption and in the share of health

consumption in augmented GDP. Nordhaus (2005) estimates that $g_h \approx 2\%$ in the later half of the 20th century, and one might take $s_h \approx .25$ as a conservative value. Table 5 considers adjustments for various values. These estimates are especially speculative given the challenges of the exercise but suggest that the social returns may go up substantially when accounting for health benefits.

The above returns credit the increase in longevity to innovations and, more particularly, to the broad range of R&D costs borne by society. Of course, several key advances in longevity have come from public health interventions, including the advance of clean water supplies in the early 20th century and anti-smoking campaigns in the later 20th century. Yet these kinds of efforts were in turn based on research insights – e.g., Pasteur’s germ theory of disease with regard to the importance of clean water, and widespread research about the harms of smoking. Cutler et al. (2006) study the historical gains in life expectancy and, considering its various potential causes, argue that “knowledge, science and technology are the keys to any coherent explanation.” At the same time, public health interventions, just like the innovations of safety glass, airbags, and other life-saving safety features in automobiles, are rooted in research and coupled with follow-on investment. This perspective further suggests that incorporating capital-investment (see Section III.A) may be appropriate for assessing social returns.

3. International spillovers

We have focused on R&D expenditure in the United States. However, the benefits of innovation in one place often spread across borders, both because ideas spillover directly and because ideas are embodied in goods and services that are traded across borders. These international spillovers mean that the innovation in the U.S. brings additional spillover benefits beyond U.S. standards of living. They also mean that some of the U.S. gains are due to innovations that come from other countries.

To examine these international spillovers, one can broaden the lens to a set of advanced economies. Here we consider the G-7 and the OECD.¹⁸ While it is difficult to say which

¹⁸ The G-7 includes Germany, Japan, Great Britain, France, Italy, Canada, and the USA. The OECD has 35 member countries, including the G-7, non-G-7 members of the European Union, Turkey, Mexico, Israel, and South Korea.

country is responsible for what portion of innovative benefits, we can look at the innovation expenditure and per-capita GDP paths in these economies collectively.

The U.S. economy invests relatively heavily in R&D (2.7% of GDP). By contrast, total R&D investment as a share of total GDP is somewhat smaller for both the G-7 countries as a whole (2.5% of total GDP since 2000) and for the OECD members as a whole (2.2% of total GDP since 2000).¹⁹ At the same time, income growth per-capita is higher in the G-7 and the OECD than in the U.S. economy. Looking since 2000, per capita income growth rates across the G-7 averaged 2.2% and across the OECD averaged 3.5%. Table 6 considers the implications for the average social returns.

Because R&D investment has declined compared to the baseline, and per-capita income growth rates have risen, the social returns to R&D investment appear higher. One may further want to include capital deepening as an important aspect in realizing these innovation gains. There is little difference, however, since 2000, in the rates of gross domestic capital formation between the U.S. and the OECD average. Thus, this additional correction, allowing for the role of capital embodiment, will be similar to that in Table 3. As in Table 3, a broader and more conservative way to capture capital deepening is to allow for economy-wide deepening that maintains the overall capital-output ratio. Because output is rising faster on average in the G-7 or OECD, this correction calls for greater net capital investment than in the U.S. economy. We also consider this correction in Table 6. Net of this greater investment cost, we still see an increase in the social returns compared to analyzing the U.S. alone.

Going further, international spillovers of innovation do not stop at the boundaries of the OECD members. The growth in standards of living in developing countries, including China and India, also benefits from innovations in advanced economies. This is true both for productivity gains and health benefits. These advances will be typically felt with a more substantial delay, as developing countries catch up later in many technologies, but with the OECD representing just 17% of the world's population, the potential scale of these spillovers is large. While we do not

¹⁹ These numbers are calculated using OECD data (<https://data.oecd.org/rd/gross-domestic-spending-on-r-d.htm>). R&D investment outside the U.S. has been increasing among the G-7 and among OECD members over time.

attempt to calculate spillovers beyond the frontier economies, such broader spillovers suggest that the international spillovers corrections in Table 6 are conservative.

Table 6: The Average Social Returns, Advanced Economies as a Whole

	R&D (%GDP)	Net Investment (%GDP)	Growth Rate per Capita (%)	Corrective Factor (β)	Average Social Benefit- Cost Ratio (ρ)	Average Social Rate of Return (r^*)
<i>Panel A: USA</i>						
Baseline	2.7	--	1.8	--	13.3	67%
Net Domestic Private Inv.	2.7	4.0	1.8	--	5.3	27%
All Capital Deepening	2.7	6.3	1.8	--	4.0	20%
<i>Panel B: G-7</i>						
Baseline	2.5	--	2.2	1.32	17.6	88%
Net Domestic Private Inv.	2.5	4.0	2.2	1.26	6.8	34%
All Capital Deepening	2.5	7.7	2.2	1.08	4.3	22%
<i>Panel C: OECD</i>						
Baseline	2.2	--	3.5	2.38	31.8	159%
Net Domestic Private Inv.	2.2	4.0	3.5	1.59	8.5	57%
All Capital Deepening	2.2	12.3	3.5	1.21	4.8	24%

Note: The corrective factor β applied here is the ratio of the social returns to the relevant row for the U.S. economy in Panel A.

C. Summary

We have examined several reasons that the baseline calculation of the average social returns to innovation investments may be too high or too low. Taking just the conservative corrections, the social benefits of innovation investments appear to exceed their costs, and substantially. Adding in natural corrections in the upward direction – due to inflation bias, health gains, or international spillovers – further increase the social benefit-to-cost ratio. Overall, it appears that a conservative estimate of the average social gains is about \$5 in benefit per \$1 invested. Considering reasonable amounts of inflation bias or health benefits can easily push the average benefit to \$10 or even \$20 per \$1 invested. These gains are just in terms of the U.S. economy. Incorporating international spillovers extends the benefits further. In sum, analyzing the average returns from a wide variety of perspectives suggests that the social returns are remarkably high.

IV. The Average Return vs. the Marginal Return

The analysis in Sections 2 and 3 are explicitly about the *average* social return to innovation investments. By focusing on the average return, we have gained several advantages. We deployed an aggregate-level analysis that can include all R&D costs, including successes and failures, and integrate across complex spillovers inherent to the innovation process. The method can also be leveraged in a transparent manner to assess forces that are not well addressed in the social returns to R&D literature. These include issues of capital embodiment, lagged effects, productivity measurement, health benefits, and international spillovers. By focusing on the average returns, we stepped past difficult conceptual and empirical issues in how to assign returns to the various components. Overall, we find that the average social returns to innovative investments appear very large.

At the same time, from a policy perspective, we may be particularly interested in the *marginal* returns. That is, we are interested not just in the average social return to what society already does, but also in the social return to undertaking additional investment in innovation. In this section, we consider the extent to which the average social returns are instructive regarding the marginal social returns. Both micro and macro considerations help inform this question, as follows.

A. Empirical Studies on the Margin

A direct, empirical argument for high marginal returns comes from existing micro literature on the social returns to formal R&D. These studies mainly investigate how variation in R&D expenditure (i.e., marginal variation) within firms and industries predicts future productivity growth of these firms and industries. These studies tend to find large social returns to increasing formal R&D, both through spillovers across firms within an industry and through technological spillovers on firms in technologically-related industries (e.g., Hall et al. 2010). While these studies use regression methods and face difficulties in interpretation, some papers advance causal identification strategies and find large social returns from marginal increases in R&D expenditure by businesses (Bloom et al. 2013).²⁰ Appendix B reviews this complementary

²⁰ See also Azoulay et al. (2019) for a causal empirical analysis that finds high returns to increased NIH funding.

literature. Overall, the social rates of return seen in these firm and industry studies, like the average returns seen in our broader calculations, show that the social returns appear very high.

B. Micro-Reasoning on the Margin

Two, related conceptual perspectives can help explain why one would expect the average social returns and marginal social returns to both be high – and even similarly high. In particular, we can write the social returns as

$$r_{social} = r_{private} + r_{spillovers}$$

for a given project.

Empirically, the calculations indicate that r_{social} appears, on average, much larger than standard private rates of return, $r_{private}$. This implies that $r_{spillovers}$ has a large positive mean. It then follows under many distributional assumptions relating $r_{private}$ and $r_{spillovers}$ that the marginal social returns will remain high. For example, if the spillover from a given project is drawn independently from the private return, or if the spillover is drawn as some multiple of the private return, then the marginal social return would remain in excess of the marginal private return. To the extent that most of the social returns are in the spillovers (as suggested by our calculations), it is then natural that the average and the marginal social returns will both be high.

A related point follows from the inherent uncertainty in innovation. As Arrow emphasized, it is typically unclear at the start whether a new R&D project is likely to bear fruit (Arrow 1962). Even when one gets close to market applications, failure is a regular occurrence. For example, most entrepreneurial ventures fail (Kerr and Nanda 2009), and those who bet on new ventures – venture capitalists – have substantial trouble predicting which investments will succeed. Kerr et al. (2014) studied a prominent venture capital firm and found that the partners' initial scoring of investment opportunities were almost entirely non-predictive of the future returns to each startup.

This uncertainty is presumably even more acute for basic research, where failure appears very common and the ultimate applications of ideas are hard to predict. For example, a basic research insight like Einstein's general relativity turned out to be an essential tool underlying the Global Positioning System (GPS), which in turn is essential to many technology applications, including new business models like Uber and Lyft (Ahmadpoor and Jones 2017). Similarly, basic research

insights into extremophile bacteria in Yellowstone National Park provided the essential gene replication technologies that underpin the biotechnology industry (Azoulay and Li, forthcoming). The connections between basic research and its ultimate applications appear broad, deep, and hard to predict (Ahmadpoor and Jones 2017, Azoulay et al. 2019).

Overall, the uncertain nature of innovation suggests that it is difficult to know the marginal return to an R&D investment, especially for basic research, applied research with uncertain endpoints, or bets on transformative business models. This uncertainty makes it difficult for investors (private or public) to credibly assess the expected returns to such innovation projects. If we are unable to predict the returns of such projects, then the marginal returns to additional investment in such projects may not be too different from the average return of the investments undertaken.

C. Macro-Reasoning on the Margin

A final and more explicit approach to assessing the marginal return to innovation is to use growth models. While the average return generalizes across a wide class of growth models, specific calculations of the marginal return will be model-dependent. Here we consider the marginal returns by considering the two primary types of innovation models that feature prominently in the endogenous growth literature.

The original approach to endogenous growth theory emphasizes that the rate of advance increases linearly in the level of R&D effort (Romer 1990, Aghion and Howitt 1992). In particular, assume for simplicity that

$$g_A = \gamma L_R \tag{10}$$

where L_R is the number of R&D workers. Intuitively, this leads to the following result, which we show formally in the Appendix.

Lemma 1: For the knowledge production function (10), the marginal social return to R&D is

$$\rho_{\text{marginal}} = \frac{g/r}{x/y} \tag{11}$$

That is, the marginal return to additional R&D effort is the same as its average return. This result is intuitive and follows because there are no diminishing returns to R&D effort. When an extra unit of effort at R&D creates a linear increase in the growth rate, the marginal return is

identical to the average return. Thus under classic, Romer-style endogenous growth models, the marginal return to additional R&D, which may be the relevant policy question, appears large and exactly in line with the calculations featured in prior sections of this paper.

At the same time, there are strong empirical arguments in the growth literature that call for an alternative model of how innovative inputs map into productivity growth. Namely, empirical evidence points to growing research effort as a requirement to drive a constant growth rate in productivity (Jones 1995, Kortum 1997, Jones 2009, Bloom et al. 2020). The macro facts suggest a relationship of the form

$$g_A = \delta A(t)^{\theta-1} L_R(t)^\sigma \quad (12)$$

which generalizes (10) on two dimensions. First, it allows for varying degrees of intertemporal spillovers: as the level of productivity, $A(t)$, advances with time, further growth may become easier or harder according to the parameter θ . Second, this generalization allows for varying degrees of diminishing returns to R&D effort at a point in time, via the parameter σ . This model produces the steady-state solution where $g_A = \frac{\sigma}{1-\theta} n$ and n is the growth rate in effort applied to R&D.²¹

This empirically-grounded knowledge production function leads to the following generalization of the marginal social return to innovation investments, which we show formally in the appendix.

Lemma 2: For the generalized knowledge production function in (12), the marginal social return to R&D is

$$\rho_{marginal} = \frac{\sigma}{1 + (\sigma - \theta)(g/r)} \frac{g/r}{x/y} \quad (13)$$

That is, the marginal return to additional R&D effort is a multiplicative factor of the average return we have calculated in Sections 2 and 3. Connecting to our prior notation, we define this multiplicative factor as $\beta_{marginal}$, where

$$\beta_{marginal} = \frac{\sigma}{1 + (\sigma - \theta)(g/r)} \quad (14)$$

²¹ The steady-state solution requires that $\theta < 1$, which means that the degree of positive intertemporal spillovers cannot be too large.

This result has several straightforward properties. First, it reproduces the Romer-style marginal return when $\sigma \rightarrow 1$ and $\theta \rightarrow 1$. Second, other things equal, increasing the degree of intertemporal spillovers (increasing θ) raises the marginal return. Third, other things equal, steeper diminishing returns to research effort (reducing σ) tends to lower the marginal return.²²

We can take one step further and calibrate the marginal returns. The marginal return depends on several empirically measured variables (g, r, x, y , which we have used to determine the average return) and two unknown parameters (σ, θ). Thus we appear to have two parameters to calibrate. However, keeping in mind that the steady-state growth rate in this model is given by $g_A = \frac{\sigma}{1-\theta}n$, we can write

$$\sigma = \frac{g}{n}(1 - \theta) \quad (15)$$

where we also observe the R&D effort growth rate n . Thus, we have in effect only one parameter to choose.²³

Table 7: The Marginal Social Returns

Intertemporal Spillovers (θ)	Implied Effort Elasticity (σ)	Marginal Return Factor ($\beta_{marginal}$)	Marginal Social Benefit-Cost Ratio ($\rho_{marginal}$)	Marginal Social Internal Rate of Return ($r_{marginal}^*$)
-0.75	1.58	0.86	11.4	101%
-0.5	1.35	0.81	10.8	87%
-0.25	1.13	0.75	10.0	73%
0	0.90	0.68	9.0	58%
0.25	0.68	0.59	7.8	45%
0.5	0.45	0.46	6.1	30%
0.75	0.23	0.28	3.7	16%

Table 7 presents calibrations of $\beta_{marginal}$ as well as the marginal social return. We use the values for g and r as in Section 2 and use $n = 2\%$, which is the growth rate of the U.S. labor

²² This result requires $\theta < 1$ and $g < r$ as conditions for the existence of steady-state growth. Note also that it is possible that the marginal return exceeds the average return. This can occur, for example, if there are increasing returns to research effort at a point in time ($\sigma > 1$) and the intertemporal spillovers are not too small.

²³ If we believe we have good information about the degree of intertemporal spillovers, θ , then we can infer σ . Or conversely we could infer θ if we believe we have evidence on σ .

force since 1960. We then consider different ranges of intertemporal spillovers from substantially negative ($\theta = -0.75$) to substantially positive ($\theta = +0.75$).

Table 7 suggests that the marginal social returns to additional R&D tend to be high. They diminish when there are steep diminishing returns to additional R&D. That is, as additional R&D effort becomes increasingly duplicative of existing R&D, or more broadly additional R&D workers at a point in time are increasingly ineffective, the marginal returns drop.²⁴

The literature does not provide clear empirical guidance on the degree of intertemporal spillovers (θ). In models where we “fish out” the pond of ideas or, more generally, the innovative search process becomes more costly as we advance, we expect $\theta < 0$ (Kortum 1997, Jones 2009). However, to the extent that new ideas or tools (e.g., calculus, computers) become fruitful inputs into innovative search, we might expect $\theta > 0$ (Weitzman 1988). Using a model of economic growth and population growth over the very long run, Kremer (1993) provides analysis that suggest a value of θ in the 0.1 to 0.4 range, which would suggest large marginal returns to additional innovative effort, but it’s not clear that such values apply in the modern era.

Embedded in the possibility of diminishing returns at the macro level are two underlying notions. First, there may be limited additional innovative lines to pursue, given the stock of current knowledge. Second, there may be limited additional innovative talent in the population. However, micro-evidence appears inconsistent with these constraints. As discussed above, studies of the returns to marginal increases in R&D by firms or in basic research at the NIH point to high marginal returns (Bloom et al. 2013; Azoulay et al. 2019). These findings are inconsistent with substantial “idea constraints” in those settings. Regarding “people constraints,” there appear to be large opportunities to expand the innovative talent pool in the US. Expansions in innovative labor could occur through immigration channels (see Kerr and Kerr, this volume) and through childhood policies (see Van Reenen, this volume). Key constraints on innovative

²⁴ An interesting feature of this calibration is that more positive intertemporal spillovers make the marginal returns to additional innovation decline. While positive intertemporal spillovers seem directly to be an advantage to the social returns, the calibration here requires that large positive intertemporal spillovers be offset by increased diminishing returns (or vice versa). This is required to match observed growth rates. An implication is that, even if one thinks there are steep diminishing returns to additional effort, the marginal social returns to additional effort can still tend to be high. For example, taking $\sigma = 0.23$ (see Table 7), one is assuming that increasing the investment in innovation by 100% would increase innovative output by only 17%. Yet the internal social rate of return is still 16% per annum (see Table 7); large, positive intertemporal spillovers makes this additional effort still highly valuable.

human capital appear through limited exposure to these career pathways, not in available talent (Bell et al. 2019a, 2019b; Azoulay et al. 2020). Immigration, education, and career-exposure policies suggest substantial short-run and long-run potential to cultivate additional innovative capacity and accelerate improvements in standards of living.

V. Conclusion

This paper considers estimates of the social returns to investments in innovation. We have introduced a transparent method that incorporates both innovative successes and failures, while also incorporating manifold externalities at work in the innovation process, including imitation, business stealing, congestion, and intertemporal spillovers. The approach can further engage a range of first-order issues that are often not considered in assessing the social returns. These dimensions include the role of capital investment, diffusion delays, productivity mismeasurement, and health outcomes among others.

Overall, we find that the average social returns to innovation investments appear very large. If formal R&D and new venture creation drive the bulk of productivity gains, then the social returns to these investments appear enormous. If a much broader set of investments, including capital embodiment, are needed to fulfill these productivity gains, then the social returns to these broader activities still appear large. Even under very conservative assumptions, it is difficult to find an average return below \$4 per \$1 spent. Accounting for health benefits, inflation bias, or international spillovers can bring the social returns to over \$20 per \$1 spent, with internal rates of return approaching 100%.

We further consider how these average returns may relate to the marginal return of additional investment in innovation. Using various perspectives, motivated by the micro and macro literatures on innovation, there are good reasons to believe that the marginal returns are also high. The implication is that policies to support further innovation investment are broadly well-motivated. Innovation investments can credibly raise economic growth rates and extend lives, paying for their costs many times over. And because the social returns exceed the private returns, public policy has a central role, and opportunity, in unleashing these gains.

References

- Adams, J.D. (1990). “Fundamental Stocks of Knowledge and Productivity Growth,” *Journal of Political Economy* 98 (3), 673–702.
- Aghion, P., and P. Howitt. (1992). “A Model of Growth through Creative Destruction,” *Econometrica* 60, 323–351.
- Ahmadpoor, M. and B.F. Jones (2017). “The Dual Frontier: Patented Inventions and Prior Scientific Advance,” *Science* 357, 583–587.
- Arrow, Kenneth (1961). “The Economic Implications of Learning by Doing,” Technical Report No. 101, Stanford Institute for Mathematical Studies in the Social Sciences.
- Arrow, Kenneth (1962). “Economic Welfare and the Allocation of Resources for Invention,” in *The Rate and Direction of Inventive Activity: Economic and Social Factors*, 609-625. Princeton, NJ: Princeton University Press.
- Azoulay, Pierre, Josh Graff Zivin, and Danielle Li (2019). “Public R&D Investments and Private-sector Patenting: Evidence from NIH Funding Rules,” *Review of Economic Studies* 86, 117-152.
- Azoulay, Pierre, Benjamin F. Jones, J. Daniel Kim, and Javier Miranda (2020). “Age and High-Growth Entrepreneurship,” *American Economic Review: Insights* 2, 65-82.
- Azoulay, Pierre, Wesley H. Greenblatt, and Misty L. Heggeness. (2020). “Long-term Effects from Early Exposure to Research: Evidence from the NIH ‘Yellow Berets’,” NBER Working Paper #26069.
- Azoulay, Pierre and Danielle Li. “Scientific Grant Funding,” in *Innovation and Public Policy*, A. Goolsbee and B. F. Jones (eds.), University of Chicago Press, forthcoming.
- Bell, Alexander, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen (2019a). “Who Becomes an Inventor in America? The Importance of Exposure to Innovation,” *Quarterly Journal of Economics* 134 (2), 647–713.
- Bell, Alexander, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen (2019b). “Do Tax Cuts Produce More Einsteins? The Impacts of Financial Incentives vs. Exposure

- to Innovation on the Supply of Inventors,” *Journal of the European Economic Association* 17 (3), 651–677.
- Bernstein, J.I., Mohnen, P. (1998). “International R&D Spillovers between U.S. and Japanese R&D Intensive Sectors,” *Journal of International Economics* 44, 315–338.
- Bernstein, J.I., Nadiri, M.I. (1989). “Research and Development and Intra-industry Spillovers: An Empirical Application of Dynamic Duality,” *Review of Economic Studies* 56, 249–269.
- Bessen, James (2015). *Learning by Doing: The Real Connection between Innovation, Wages, and Wealth*, Yale University Press: New Haven.
- Bloom, Nicholas, Mark Schankerman, and John Van Reenen (2013). “Identifying Technology Spillovers and Product Market Rivalry,” *Econometrica* 81, 1347-1393.
- Bloom, Nicholas, Charles Jones, John Van Reenen, and Michael Webb. (2020). “Are Ideas Getting Harder to Find?” *American Economic Review* 110, 1104-44.
- Bresnahan, Tim (1986). “Measuring Spillovers from “Technical Advance”,” *American Economic Review* 76, 741–755.
- Brouwer, E. and K. Kleinknecht (1997). “Measuring the Unmeasurable: A country's Non-R & D Expenditure on Product and Service Innovation,” *Research Policy* 25, 1235-1242.
- Brynjolfsson, E., A. Collis, and F. Eggers (2019). “Using Massive Online Choice Experiments to Measure Changes in Well-being,” *Proceedings of the National Academy of Sciences* 116, 7250-7255.
- Coe, D.T., Helpman, E. (1995). “International R&D Spillovers,” *European Economic Review* 39, 859–887.
- Cutler, D., A. Deaton, and A. Lleras-Muney. (2006) “The Determinants of Mortality,” *Journal of Economic Perspectives* 20, 97-120.
- Evangelista, E., T. Sandven, G. Sirilli, and K. Smith (2010). “Measuring Innovation in European Industry,” *International Journal of the Economics of Business* 5, 311-333.

- Evenson, R. (2001). "Economic impacts of agricultural research and extension," in the *Handbook of Agricultural Economics*, 573-628.
- Goolsbee, Austan and Peter Klenow. (2018). "Internet Rising, Prices Falling: Measuring Inflation in a World of E-Commerce," *American Economic Review* 108, 488-492.
- Gordon, Robert (1999). "The Boskin Commission Report and its Aftermath," *Monetary and Economic Studies*, Institute for Monetary and Economic Studies, Bank of Japan, vol. 17, 41-68.
- Griffith, R., Redding, S., Van Reenen, J. (2004). "Mapping the Two Faces of R&D: Productivity Growth in a Panel of OECD Manufacturing Industries," *Review of Economics and Statistics* 86 (4), 883–895.
- Griliches, Zvi (1958). "Research Cost and Social Returns: Hybrid Corn and Related Innovations," *Journal of Political Economy* 66, 419–431.
- Grossman, V., T. Steger, and T. Trimborn. (2016). "Quantifying Optimal Growth Policy," *Journal of Public Economic Theory* 18, 451-485.
- Guzman, Jorge and Scott Stern (2017), "Nowcasting and Placecasting Entrepreneurial Quality and Performance" Chapter in NBER book *Measuring Entrepreneurial Businesses: Current Knowledge and Challenges*, pages 11-62 National Bureau of Economic Research, Inc., John Haltiwanger, Erik Hurst, Javier Miranda, and Antoinette Schoar, editors.
- Hall, Bronwyn, Mairesse, Jacques, and Pierre Mohnen (2010). "Measuring the Returns to R&D," in *Handbook of the Economics of Innovation*, 1033-1082.
- Jones, Benjamin F. (2009). "The burden of knowledge and the 'death of the renaissance man': Is innovation getting harder?" *Review of Economic Studies* 7, 283–317.
- Jones, Charles (1995). "R&D-based Models of Economic Growth," *Journal of Political Economy* 103, 759–784.
- Jones, Charles and John Williams (1998). "Measuring the Social Rate of Return to R&D," *Quarterly Journal of Economics* 113 (4), 119–135.

- Kao, C., Chiang, M., Chen, B. (1999). “International R&D Spillovers: An Application of Estimation and Inference in Panel Cointegration,” *Oxford Bulletin of Economics and Statistics* 61, 691–709.
- Kerr, William and Ramana Nanda (2009). “Democratizing Entry: Banking Deregulations, Financing Constraints, and Entrepreneurship,” *Journal of Financial Economics* 94(1): 124–49.
- Kerr, William, Ramana Nanda, and Matthew Rhodes-Kropf (2014). “Entrepreneurship as experimentation,” *Journal of Economic Perspectives* 28, 25-48.
- Kerr, Sari and William Kerr. “Immigration Policy Levers for US Innovation and Startups,” in *Innovation and Public Policy*, A. Goolsbee and B. F. Jones (eds.), University of Chicago Press, forthcoming.
- Kortum, Sam (1997). “Research, Patenting, and Technological Change,” *Econometrica* 65, 1389–1419.
- Leonard, W.N. (1971). “Research and Development in Industrial Growth,” *Journal of Political Economy* 79 (2), 232–256.
- Mansfield, E., Rapoport, J., Romeo, A., Wagner, S., Beardsley, G. (1977). “Social and Private Rates of Return from Industrial Innovations,” *Quarterly Journal of Economics* 77, 221–240.
- Mokyr, Joel (1990). *The Lever of Riches: Technological Creativity and Economic Progress*. Oxford University Press: New York.
- Murphy, Kevin M. and Robert Topel (2007). “The Value of Health and Longevity,” *Journal of Political Economy* 114, 871-904.
- Nadiri, M.I., Kim, S. (1996b). “International R&D Spillovers, Trade and Productivity in Major OECD Countries,” NBER Working Paper 5801.
- Nordhaus, William. (2005). “Irving Fisher and the Contribution of Improved Longevity to Living Standards,” *American Journal of Economics and Sociology* 64, 367-392.

- Rachel, Lukasz and Lawrence H. Summers. (2019). "On secular stagnation in the industrialized world," NBER Working Paper No. 26198.
- Romer, Paul M. (1990). "Endogenous Technological Change," *Journal of Political Economy* 98, S71–S102.
- Solow, Robert M. (1956). "A Contribution to the Theory of Economic Growth," *Quarterly Journal of Economics* 70, No. 1, 65-94.
- Tewksbury, J.G., Crandall, M.S., Crane, W.E. (1980). "Measuring the Societal Benefits of Innovation," *Science* 209, 658–662.
- Trajtenberg, Manuel (1989). "The Welfare Analysis of Product Innovations, with an Application to Computed Tomography Scanners," *Journal of Political Economy* 97 (2), 444–479.
- Van Reenen, John. "Innovation and Human Capital Policy," in *Innovation and Public Policy*, A. Goolsbee and B. F. Jones (eds.), University of Chicago Press, forthcoming.
- Von Hippel, Eric. (1976). "The Dominant Role of Users in the Scientific Instrument Innovation Process," *Research Policy* 5, 212-239.
- Westmore, Ben (2013). "R&D, Patenting and Growth: The Role of Public Policy," OECD Economics Department Working Papers No. 1047.
- Weitzman, Martin L. (1998). "Recombinant Growth," *Quarterly Journal of Economics* 113, 331–360.
- Wolff, E.N., Nadiri, M.I. (1993). "Spillover Effects, Linkage Structure, and Research and Development," *Structural Change and Economic Dynamics* 4 (2), 315–331.
- Wright, T.P. (1936). "Factors Affecting the Cost of Airplanes," *Journal of the Aeronautical Sciences* 3, 122-128.

Appendix A: Spillovers and the Social Returns to Innovation

The society-wide gains from an innovation may differ substantially from the private gains to the innovator. The difference comes from the many potential “spillovers” that can follow from the creation and introduction of new ideas. In this appendix, we discuss the range of these potential spillovers, which in turn make the measurement of the social return to innovation challenging.

Imitative Spillovers

The innovative investments made by one firm may not only raise the investing firm’s productivity, but also raise the productivity of other firms. In particular, other firms can imitate the advance (e.g., Segerstrom 1991). For example, consider computer manufacturers. When a more advanced microprocessor, memory chip, or monitor is created, competing firms will see and learn from these innovations and improve their own products. These “imitative” knowledge spillovers increase the social returns to innovation, even as imitation by competitors may reduce the private return to the original innovator. Beyond product innovations, process innovations – such as Henry Ford’s assembly line, Geoff Hinton’s artificial intelligence algorithms, or the WHO’s surgical checklist – can also be learned and imitated by others, extending the benefits far beyond the original innovator.

User Spillovers

An important, second potential spillover is the benefit that accrues to users (e.g., Trajtenberg 1989). For example, more advanced computing machinery will presumably increase the productivity of the downstream firms that purchase and deploy these more advanced machines. This user benefit is not likely to be fully captured by the upstream innovator: in particular, the users who buy the product presumably expect a benefit in excess of the product’s price. User spillovers can occur between firms in a vertical supply relationship. They can also occur for the end user – the consumer – creating consumer surplus that is not captured by the innovating firm. The downstream benefits may be especially large when the upstream producers are competitive and imitate each other’s innovations (e.g., Petrin 2002).

Intertemporal Spillovers

A potentially central but difficult to estimate spillover is intertemporal in nature, where a given advance may influence the capacity for future advances (e.g., Romer 1990, Scotchmer 1991, Weitzman 1998). This intertemporal element could involve opening research avenues in a given product line: a specific advance in jet engine design may inspire a stream of future jet engine innovations. The spillovers may also be far more general. For example, technologies like electricity, computers, and mobile phones serve as platforms for enormous arrays of future innovations. Taking smartphones as one example, these tools have spurred the innovation of millions of new software applications.²⁵ Mobile phones have also sparked the creation of transformative business models, including mobile payments and the ride-sharing industry.

When the intertemporal spillovers on future innovation are broad, it becomes difficult to measure the social returns to the original innovation. With general purpose technologies like mobile phones, the Internet, computers, lasers, or electricity, it is difficult even to enumerate the full set of future applications that build upon them. A question like “what is the social return to the Internet?” is difficult to answer because the applications are so varied.

This difficulty is also acute with basic research. By definition, basic research is not directed at specific marketplace innovations. Rather, it is intended to advance understanding and introduce new ideas upon which future applications may build. Essentially, the marketplace returns to basic research are *all* in the intertemporal spillovers. And while basic research is an uncertain exercise riddled with failure, it also produces insights that are ultimately essential to marketplace innovations and socioeconomic prosperity. For example, without basic research breakthroughs in genetics – from Mendelian inheritance to Watson and Crick’s structure of DNA to Kary Mullis’s polymerase chain reaction – there would be no biotechnology industry and many of our most advanced medical treatments would not exist. Advances in mathematics, chemistry, solid-state physics, material science, and statistics, to name just a few fields, underpin substantial marketplace applications (Ahmadpoor and Jones 2017). Asking “what is the social return to

²⁵ In 2019, there are 2.47 million apps available on the Android platform and 1.80 million apps available on the Apple platform.

learning the structure of DNA?” or “what is the social return to calculus?” are obviously difficult, because, once again, the applications are so varied.

The above discussion suggests that the intertemporal spillovers are largely positive, as an advance can facilitate future advances. But it is also possible that intertemporal spillovers are negative. The main reason for a negative intertemporal spillover is that we may find discovery of new ideas increasingly difficult. For example, if ideas are fruit on a tree, we may naturally pluck the low-hanging fruit first. Then future innovation will become harder to achieve. There is substantial micro and macro evidence along these lines, where innovation requires more effort with time (C. Jones 1995; B. Jones 2009; Bloom et al. forthcoming). But it remains an open question whether the intertemporal spillovers are on net positive or negative.

Business Stealing

Returning to firms, additional issues could limit the social return to innovation. In particular, in a competitive context it is possible that the social returns are actually below the private return. This effect comes from “business stealing”, where the advance of one firm may come in part by stealing business from other firms (e.g., Aghion and Howitt 1998). Concretely, consider a small innovation that allows a firm to produce a piece of machinery at a slightly lower cost than all the other firms in a competitive market. This innovating firm may then grow to take over the market and see an enormous private return, but the social return may actually be very small. More generally, any time a firm or industry is growing at the expense of other firms and industries, looking narrowly at the private R&D returns to the advancing firm or industry will tend, other things equal, towards overstating the social returns.

Duplication

A final kind of negative spillover comes within the R&D process itself, when research teams duplicate each other’s efforts (e.g., Dixit 1998). For example, many firms may simultaneously seek to create the same new technology. Similarly, multiple basic research teams may race toward the same experimental result. Because research teams do not internalize their effect on the other teams, there may be too much entry on a given research line.

Appendix B: Empirical Estimates of the Social Returns to R&D: Existing Literature

This appendix reviews existing approaches to calculating the social returns to R&D. We review technology case studies, firm and industry level studies, and country-level studies. This literature use a variety of methods and provides a series of informative findings. The typical finding is that the social returns appear very large. At the same time, each approach has methodological limitations.

Technology case studies

The “case study” approach compares R&D costs with their associated benefits for specific technologies and sectors. Griliches (1958), in a seminal contribution, considered the development of hybrid corn in the United States. The R&D costs targeting hybrid corn are counted up over several decades. The benefits are calculated as the increased corn output that results, net of increased input costs. Both the R&D costs and the production benefits are summed up at a point in time using an assumed discount rate. In Griliches’s central estimate, the social returns appear very large: 1\$ of R&D costs provides a \$7 net present value benefit.²⁶

Other case studies have examined numerous agriculture innovations (e.g., Evenson 2001) and small sets of industrial innovations – including mechanical, chemical, electronics, and consumer product innovations (Mansfield et al. 1977, Tewksbury et al. 1980). Bresnahan (1986) studied mainframe computers in financial services. Trajtenberg (1989) studied CT scanners and their benefits for health care. While estimates vary, these studies typically show large social returns. For example, studies of public agricultural research suggest social returns typically above 40% (Evenson 2001). The findings across 37 industrial innovations studied in Mansfield et al. (1977) and Tewksbury et al. (1980) suggest a median social rate of return of 71%.

The primary challenge for case studies is whether they generalize. Hybrid corn, mainframe computers, and CT scanners were successful innovations. Case studies of failures are rare, even though failure in innovation is common (Arrow 1960; Kerr et al. 2014). By leaving out failures, case studies may overstate the general social returns to R&D. On the other hand, case studies

²⁶ Griliches argues this estimate is conservative. He uses a high discount rate (10%) and other conservative assumptions to argue that the social returns are at least \$7 per \$1 spent. This equates to an internal rate of return of at least 35-40%.

focus on narrow innovations or applications. The social returns for far-reaching innovations -- electricity, lasers, computers, and gene sequencing -- are very hard to calculate and yet may have the highest social returns of all. Thus, whether the case study evidence overstates or understates the average social returns is unclear.

Firm and industry analyses

A separate literature uses regression methods to study the social returns to R&D. In these regressions, the dependent variable is typically the output or productivity of the firm or industry. The explanatory variables are R&D expenditures. At the firm level, a private return is estimated by looking at how a firm's own R&D expenditure predicts that firm's output or productivity growth. Social returns are incorporated by further examining how the focal firm or industry's output growth depends on the R&D investments by *other* firms or industries. This cross-firm or cross-industry spillover is estimated in the regression context by including "outside R&D" as a separate predictor of a given firm or industry's outcome.

Regression approaches often find substantial social returns. Hall et al. (2010) review the regression evidence and suggest that private returns to R&D are most likely in the 20-30% range. Estimates of the cross-firm or industry spillovers tend to be additionally positive, but these estimates vary considerably across studies and are often imprecise. For example, some studies have suggested large returns can be captured from outsider R&D (Griliches and Lichtenberg 1984; Griffith et al. 2004), while others have suggest that there can be little or no return from outsider R&D (Bernstein and Nadiri 1989; Wolff and Nadiri 1993; Bernstein and Mohnen 1998).

The regression approach embodies a number of assumptions. First, to interpret a regression coefficient as a rate of return, one must assume specific production functions relating R&D to productivity growth. Second, one must make assumptions about lags, since the output growth today may depend not just on last year's R&D expenditure, but on R&D projects begun in prior years. In practice, regression methods typically assume a very rapid payoff of R&D. Third, one must make assumptions about the scope of spillovers, where firms or industries that are nearer in technology may have more spillover potential. Hall et al. (2010) review the various assumptions authors have used on these dimensions, and the variety of assumptions may help explain variant results.

The regression results also do not imply causation. A positive correlation between R&D expenditures and firm output could be due to reverse causation or omitted variables. Firms with high output growth may choose to do substantial R&D, so that the causation runs backwards. And good technology prospects may cause all firms to do more R&D and also see output increases; the apparent spillover from “outside R&D” may then be a spurious association driven by common technology opportunities. Thus interpreting private or social returns from simple regressions is not straightforward.

In light of these issues, two studies are notable for attempting to causatively estimate the social returns. Bloom et al. (2013) use changes in federal and state-level R&D tax incentives, which change the R&D costs of firms. These authors show that R&D expenditures go up when a firm’s tax costs go down. The resulting change in R&D investment in turn drives greater firm growth and greater spillovers on other firms. Bloom et al. (2013) estimate a private return of 21% and a social return of 55%.²⁷

A limitation for all regression models of spillovers is that they must take a stand on the boundary of spillovers. Anything “further outside” the outside R&D measures is, by construction, ignored. One omission is basic research, including that conducted in universities and government laboratories. These R&D investments are left out of industry studies but may have important effects. Azoulay et al. (2019) tackle basic research spillovers in the context of biomedical innovation. Using shocks to National Institutes of Health funding allows this study to make a causative interpretation. They find that an additional \$10 million dollars in NIH funding leads to 2.7 additional private sector patents. By imputing market values to these patents, the authors calculate a commercial return of at least \$2 per \$1 spent by the NIH. The social returns, which would require assessing the net health advantages of these innovations, would presumably be higher. An additional important finding is that half of the patents come in disease areas outside the target of the NIH funding, which points to the broad scope of basic research spillovers.

²⁷ An important feature of this study is that it confronts two dimensions of R&D spillovers. The effect of ‘outside R&D’ is in principal a mixture of two forces above. First, there may be knowledge spillovers, where technology advances in one firm are absorbed by other firms, raising these other firms’ productivity. Second, there may be business stealing, where the gains by one firm may come at the expense of other firms’ business. Bloom et al. (2013) distinguish these channels by separately considering firms that are close together in technology space (allowing knowledge spillovers) and firms that are close together in product space (allowing business stealing). The finding that the social returns are 55% nets out both channels, suggesting that knowledge spillovers dominate.

Country-level analyses

Regression models can also be conducted at the national level. Here the dependent variable is national total factor productivity, and the R&D input is the total R&D expenditure within the country. Cross-country spillovers are examined by including other countries' R&D as a separate explanatory variable. Coe and Helpman (1995) study 22 high-income countries. They find that a strong positive relationship between R&D expenditure and productivity growth at the national level. Taking the association as causative, the own rate of return to R&D averaged 123% in G-7 countries and averaged 85% in 15 other high-income countries. Cross-country spillovers also appear substantial, adding another 30% to the returns. Several studies consider alternative regression specifications to Coe and Helpman (1995) and alternative national-scale settings or sets of countries. The nation's returns to R&D always appear positive in these studies, but the magnitudes vary considerably, with some studies finding very large returns (e.g., Kao et al. 1999) and others finding more modest returns (Nadiri and Kim 1996, Westmore 2013).

An important advantage of aggregation to the national level is that it can include all R&D (including basic research expenditure) and net out R&D spillovers across firms and industries, including knowledge spillovers and business stealing effects. The "own return" in a national regression is thus conceptually much closer to a social return than the narrower technology, firm, or industry level analyses. The cross-country aspect of Coe and Helpman (1995) and ensuing studies adds a further dimension of spillover, where the benefits of innovations extend beyond national boundaries.

The disadvantage of country-level regression approaches are similar to above, especially with regard to causative identification. One may be particularly concerned at the national level about spurious associations that disrupt interpretation. For example, R&D investment is responsive to business cycles, leading to reverse causation problems (Ouyang 2011, Aghion et al. 2012). More broadly, omitted variables may bias the correlations.

An alternative macroeconomic approach is model driven. Here authors use specific growth models to calculate the marginal return of additional R&D spending. Jones and Williams (1998) take this approach and show under fairly broad theoretical conditions that private returns to R&D seen in the micro literature will tend to understate the social returns. They conclude that optimal R&D investment is 2-4 times greater than observed investment. Many ensuing studies build

specific endogenous growth models and calibrate them to micro and macro evidence (e.g., Jones and Williams 2000, Sener 2008, Grossman et al. 2016). This work arrives at the similar broad conclusion where the marginal social returns to additional R&D are high and that advanced economies underinvest in R&D.

Summary

Using different methods and data, the existing literature suggests that the social returns to R&D are high. At the same time, the scope of spillovers considered is often limited, especially with studies focused on specific technologies, firms, and industries. Intertemporal spillovers, which may play out in diffuse ways and with long delays, are typically ignored. And each of the above methods has specific limitations. Despite these differences and limitations, the diversity of approaches seen across the literature reach similar conclusions: the social returns appear very high. The complementary calculation in this paper, which addresses several limitations, further indicates that high social returns to innovation investments appear robust.

Appendix C: Formal Results

A. The Baseline Average Social Returns

Here we derive the baseline estimate for the social rate of return to aggregate investment in innovation. The average social returns are calculated by comparing the observed growth path with the counterfactual growth path that would emerge in the absence of innovation investments. We will look at the benefits in per-capita income terms. For the observed growth case, we see the path of GDP per capita, $y(t)$, and the path of innovation investment per capita, $x(t)$. For the counterfactual case, we have an alternate path of GDP per capita, $\hat{y}(t)$, given an alternate path of investment, $\hat{x}(t)$.

The social returns, $\rho(t)$, are calculated by comparing the ratio of the net present value of the benefits, $B(t)$, to the net present value of the costs, $C(t)$. Namely,

$$\rho(t) = \frac{B(t)}{C(t)}$$

Let's say the counterfactual path begins at some time t_0 . Then, comparing the observed and counterfactual investment paths, the present value of the innovation benefits is

$$B(t_0) = \int_{t_0}^{\infty} [y(t) - \hat{y}(t)]e^{-r(t-t_0)} dt$$

and the present value of innovation costs is

$$C(t_0) = \int_{t_0}^{\infty} [x(t) - \hat{x}(t)]e^{-r(t-t_0)} dt$$

To proceed to our baseline estimate, we first must define a counterfactual of interest. In particular, we want to consider the average return to all innovation investment, aggregating across the many spillovers associated with this process. As a thought experiment, we can do this by “turning off” innovation at time t_0 . Thus, by definition, we write $\hat{x}(t) = x(t)$ for $t < t_0$, and we write $\hat{x}(t) = 0$ for $t \geq t_0$.

The remaining question concerns the counterfactual path of income per capita. For a simple baseline, we write $\hat{y}(t) = y(t)$ for $t < t_0$ and we write $\hat{y}(t) = y(t_0)$ for $t \geq t_0$. That is, we assume that per-capita income stops growing in the absence of further innovation investments. This baseline counterfactual path embeds a set of assumptions, and relaxing those assumptions is the subject of Section IV. But note that, while simplistic, this counterfactual path is broadly consistent with neoclassical growth theory, where (1) following Solow, growth in per-capita income requires gains in productivity and (2) following endogenous growth theory, gains in productivity come from explicit investments in innovation.

For simplicity, take the stylized facts of a balanced growth path, where the observed path of $y(t)$ grows at a constant rate g and measured innovation investment (i.e., R&D) is an approximately constant share of GDP, and thus also grows at rate g . The present value of the benefits of innovation are then

$$B(t_0) = y(t_0) \left[\frac{1}{r-g} - \frac{1}{r} \right]$$

and the present value of the costs of innovation are

$$C(t_0) = x(t_0) \left[\frac{1}{r - g} \right]$$

It then follows that the social benefit cost ratio (the amount of benefit per unit of cost) is

$$\rho = \frac{g/r}{x/y}$$

where we have dropped the time notation, t_0 , given that we are looking at a balanced growth path, where $x(t)/y(t)$ is constant. Alternatively one can describe the social rate of return, r^* . This is the discount rate for which the benefits would equal the costs ($\rho = 1$). That is,

$$r^* = \frac{gy}{x}.$$

A.1 Discrete Time Analog

As an alternative derivation, we can consider a discrete time analog. Here innovation is not “turned off” forever but rather for just one period. This approach may better clarify that our counterfactual path doesn't change the intertemporal spillovers from innovation, as this counterfactual preserves the exact same path of productivity gains, but with a one period delay.

In particular, let there be a series of investments, x_t , that improve productivity, A_t . As a counterfactual, we imagine that in some year t_0 , no such investments are made and thereafter exactly the same investments are made as along on the observed path, only one period later. That is, we consider the innovation investment path where $\hat{x}_t = x_t$ for $t < t_0$; $\hat{x}_t = 0$ for $t = t_0$; and $\hat{x}_t = x_{t-1}$ for $t > t_0$. Since these are truly identical investments (i.e., the same innovation projects), we imagine that they must have the same ultimate effect on productivity. Thus it must be, ultimately, that $A_t = A_{t-1}$. Now, in neoclassical growth theory, we have y_t/A_t equal to a constant. This implies that, ultimately, $y_t = y_{t-1}$.

In a simple, "immediate innovation effect" model, we have $\hat{y}_t = y_t$ for $t \leq t_0$; and $\hat{y}_t = y_{t-1}$ for $t > t_0$. We take this as our baseline counterfactual case, commensurate with the baseline approach in Section 2.

What are the social returns? The net present value of the difference in investment costs along the observed path and counterfactual paths is

$$C_{t_0} = x_{t_0} \left[\frac{r}{r - g} \right]$$

And the net present value of the difference in the benefits along the observed and counterfactual growth paths is

$$B_{t_0} = y_{t_0} \left[\frac{g}{r - g} \right]$$

Then, along a balanced growth path we have, once again,

$$\rho = \frac{g/r}{x/y}$$

B. Social Returns with Health Benefits

To incorporate health benefits into the social returns, we first expand the definition of GDP to include a “health consumption” component, which represents the flow value of being alive. Denote this health consumption flow as h , and define “augmented GDP per capita” as y^* , which includes this health consumption. That is,

$$y^* = y + h$$

Similarly, denote the growth rate of augmented GDP per capita as g^* . Based on the above definition of augmented GDP per capita, it follows that

$$g^* = g(1 - s_h) + g_h s_h$$

where $s_h = h/y^*$ is the share of health consumption in augmented GDP and g_h is the growth rate of h .

The true social returns to innovation will then be

$$\rho^* = \frac{g^*/r}{x/y^*}$$

which makes two adjustments compared to the baseline calculation of the social returns to innovation. First, the relevant benefit measure is based on g^* , which incorporate progress in health. Second, the relevant cost measure is still total innovation expenditure, x , but it now viewed as a share of the augmented GDP per capita measure, y^* .

Using the expressions for y^* and g^* , the health-augmented social rate of return to innovation can be written as

$$\rho^* = \left(1 + \frac{s_h}{1 - s_h} \frac{g_h}{g}\right) \frac{g/r}{x/y}$$

as presented in the text.

C. Proof of Lemmas

Lemma 1: For the knowledge production function (10), the marginal social rate of return to R&D is $\rho_{\text{marginal}} = \frac{g/r}{x/y}$.

Proof:

The output path of the economy is $Y(t) = A(t)L_Y(t)$, with workers paid a competitive wage $w(t) = A(t)$. The R&D expenditure path is $X(t) = w(t)L_R(t)$. In per-capita terms, income per capita is $y(t) = A(t)$ and R&D expenditure per capita is $x(t) = A(t) \frac{L_R(t)}{L(t)}$.

We compare the observed balanced growth path with a counterfactual path in which R&D expenditure per capita is raised by v percent. Comparing the observed income path, $y(t)$, and the counterfactual income path, $\hat{y}(t)$, the net present value of the benefits of increasing innovation investment is

$$B(t_0) = \int_{t_0}^{\infty} (\hat{y}(t) - y(t)) e^{-r(t-t_0)} dt = \int_{t_0}^{\infty} (\hat{A}(t) - A(t)) e^{-r(t-t_0)} dt \quad (16)$$

where the counterfactual path begins at time t_0 . Comparing the observed innovation expenditure path, $x(t)$, with the counterfactual innovation investment path, $\hat{x}(t)$, the net present value of the costs of increasing innovation investment is

$$\begin{aligned} C(t_0) &= \int_{t_0}^{\infty} (\hat{x}(t) - x(t)) e^{-r(t-t_0)} dt \\ &= (x/y) \int_{t_0}^{\infty} (\hat{A}(t)(1+v) - A(t)) e^{-r(t-t_0)} dt \end{aligned} \quad (17)$$

where the resource allocation, $\frac{L_R(t)}{L(t)} = x/y$, is a constant on the observed balanced growth path and is a constant that is proportionally $1 + v$ higher on the counterfactual growth path (which is also balanced in this case).

To consider the social returns to R&D, we can then integrate these expressions. Using the Romer-style knowledge production function, (10), we have²⁸

$$A(t) = A(t_0)e^{\gamma L_R(t-t_0)}$$

$$\hat{A}(t) = A(t_0)e^{\gamma(1+v)L_R(t-t_0)}$$

for the observed and counterfactual paths of productivity. The net benefits from increased innovation investment are then

$$B(t_0) = \frac{A(t_0)}{r - \gamma L_R} \left[\frac{v\gamma L_R}{r - (1+v)\gamma L_R} \right]$$

and the net costs from increased innovation investment are

$$C(t_0) = \frac{1}{x/y} \left[\frac{A(t_0)}{r - \gamma L_R} \right] \left[\frac{vr}{r - (1+v)\gamma L_R} \right]$$

The social return to any adjustment of size v is then

$$\rho_v = \frac{B(t_0)}{C(t_0)} = \frac{\gamma L_R/r}{x/y}$$

where we note that the steady state growth rate on the observed path is $g = \gamma L_R$. Thus we have

$$\rho_{marginal} = \frac{g/r}{x/y}$$

which was to be shown.

Lemma 2: For the generalized knowledge production function in (12), the marginal social return

to R&D is $\rho_{marginal} = \frac{\sigma}{1+(\sigma-\theta)(g/r)} \frac{g/r}{x/y}$.

Proof:

²⁸ Recall that the Romer-style growth models require constant population for a balanced growth path.

Using the same approach as in Lemma 1 will not work here, because in general the counterfactual path $\hat{A}(t)$ is not simply a constant, proportional change in the growth rate, as in the Romer model. However, the counterfactual path still has a closed form solution. In particular, we now have the generalized knowledge production function (12)

$$\dot{A}(t) = \delta A(t)^\theta L_R(t)^\sigma \quad (12)$$

This knowledge production function is a separable, non-linear differential equation. Separating and integrating both sides, we have solutions of the form

$$\int_{-\infty}^t \hat{A}(\tau)^{-\theta} d\hat{A}(\tau) = \int_{-\infty}^{t_0} \delta \hat{L}_R(\tau)^\sigma d\tau + \int_{t_0}^t \delta \hat{L}_R(\tau)^\sigma d\tau \quad (18)$$

On the counterfactual path, the number of R&D workers follows

$$\hat{L}_R(t) = \begin{cases} L_R(t), & t < t_0 \\ (1+v)L_R(t), & t \geq t_0 \end{cases}$$

Where $L_R(t) = L_R(t_0)e^{n(t-t_0)}$ grows at a constant exponential rate n . We therefore integrate (18) and solve for the counterfactual productivity path as

$$\hat{A}(t) = A(t_0) \left[1 + (1+v)^\sigma (e^{\sigma n(t-t_0)} - 1) \right]^{\frac{1}{1-\theta}} \quad (19)$$

where $A(t_0) = \hat{A}(t_0) = \left[\frac{(1-\theta)\delta L_R(t_0)^\sigma}{\sigma n} \right]^{\frac{1}{1-\theta}}$.

The path $\hat{A}(t)$ cannot be integrated easily into a net present value. However, one can still produce an analytic solution for the marginal social return as follows. First, write the social return as

$$\rho_v = \frac{B(t_0)}{C(t_0)} = \frac{\int_{t_0}^{\infty} (\hat{A}(t) - A(t)) e^{-r(t-t_0)} dt}{\left(\frac{x}{y}\right) \int_{t_0}^{\infty} (\hat{A}(t)(1+v) - A(t)) e^{-r(t-t_0)} dt} = \frac{1}{x/y} \left[\frac{1}{1+Q(v)} \right]$$

where

$$Q(v) = \frac{\int_{t_0}^{\infty} v \hat{A}(t) e^{-r(t-t_0)} dt}{\int_{t_0}^{\infty} (\hat{A}(t) - A(t)) e^{-r(t-t_0)} dt}$$

We seek the marginal return, where v is small. While the limit $\lim_{v \rightarrow 0} Q(v)$ is not defined in the above form, we can instead use L'Hopital's rule to write

$$\lim_{v \rightarrow 0} Q(v) = \lim_{v \rightarrow 0} \frac{\int_{t_0}^{\infty} \hat{A}(t) e^{-r(t-t_0)} dt + \int_{t_0}^{\infty} v \frac{\partial \hat{A}(t)}{\partial v} e^{-r(t-t_0)} dt}{\int_{t_0}^{\infty} \frac{\partial \hat{A}(t)}{\partial v} e^{-r(t-t_0)} dt} \quad (20)$$

The derivative of the path $\hat{A}(t)$ with respect to v , using (19), is

$$\frac{\partial \hat{A}(t)}{\partial v} = \frac{A(t_0)}{1 - \theta} \sigma (1 + v)^{\sigma-1} (e^{\sigma n(t-t_0)} - 1) [\hat{A}(t)/A(t_0)]^{\theta}$$

We can then integrate out the expressions in (20), noting that $\lim_{v \rightarrow 0} \hat{A}(t) = A(t)$, and take the limit of $Q(v)$ as

$$\lim_{v \rightarrow 0} Q(v) = \frac{r - \theta g}{\sigma g}$$

With some algebra, we can thus write

$$\rho_{marginal} = \frac{\sigma}{1 + (\sigma - \theta)(g/r)} \frac{g/r}{x/y}$$

as was to be shown.