Policies to Stimulate Innovation

How effective are policies to encourage investment in innovation by firms, and what impact do they have on the macroeconomy?

Andrew Atkeson
UCLA and Federal Reserve Bank of Minneapolis

Ariel T. Burstein
UCLA

Economic Policy Papers *are based on policy-oriented research by Minneapolis Fed staff and consultants. The papers are an occasional series for a general audience. Views expressed are those of the authors, not necessarily of others in the Federal Reserve System.*

ABSTRACT

National policymakers have long been interested in technological innovation by firms and its potential contribution to economic growth and improved well-being. Policies to encourage innovation by firms include government funding for research and development, direct and indirect subsidies, tax credits and other tax benefits such as deductibility of research expenses. Other policies such as the corporate profits tax also impact firms' decisions to innovate. Which policies are most successful in spurring innovation at companies, given their fiscal cost to taxpayers? To what extent does the firm-level innovation induced by these policies truly generate broader economic growth?

This policy paper seeks to provide insight into key considerations in innovation policy. The overarching issue is: How do policies that affect firms' innovation costs and benefits impact aggregate innovation activity, output, productivity and welfare?

We establish a benchmark model of innovation that provides a straightforward procedure for estimating relative magnitudes of long-run macroeconomic impact of a range of innovation policy options. The procedure gauges approximate impact of two innovation policies on macroeconomic outcomes quite simply, through computing and comparing the government's fiscal expenditure on these two policies. Two innovation policies have approximately the same impact on aggregate innovation, output and productivity in the long run if they have the same fiscal impact on taxpayers.

The response of economic welfare and GDP over the long run to changes in innovation policy is highly sensitive to the size of innovation spillovers; welfare gains

Continued on page 6

Introduction[1]

National policymakers have long been interested in technological innovation and its potential contribution to economic growth and improved well-being. The Obama administration has embraced innovation as "the foundation of American economic growth and national competitiveness." In launching the "Strategy for American Innovation" in November 2010, the president remarked, "[T]he key to our prosperity … as it has always been—will be to compete by developing new products, by generating new industries, by maintaining our role as the world's engine of scientific discovery and technological innovation."[2]

Policies to encourage innovation by firms include government funding for research and development (R&D), direct and indirect subsidies, tax credits and other tax benefits, such as deductibility of research expenses. Other policies not typically thought of as aimed at stimulating innovation, such as the corporate profits tax, also impact firms' decisions to innovate. But to channel support effectively, policymakers need to know which policies are most successful in spurring innovation at companies, given their fiscal cost to taxpayers, and to what extent the firm-level innovation induced by these policies truly generates broader economic growth. Also important to policy: What factors influence the effectiveness of innovation subsidies in promoting economic well-being over the long term?

In particular, the idea that innovative activity by firms has "spillovers" that promote the wider diffusion

Innovation

How do policies that affect firms' innovation costs and benefits impact aggregate output, productivity and welfare?

Andrew Atkeson

Ariel T. Burstein

ABSTRACT from page 4

could vary between virtually no change and a 50 percent increase in equivalent consumption, depending on spillover size.

Unfortunately, we cannot accurately measure these long-run effects without accurate estimates as to the magnitude of innovation spillovers. Results from our model indicate, however, that even under ideal conditions, it should be very difficult to measure spillovers using data on medium-term response of the macroeconomy from changes in innovation policy. That is, evidence from the medium term is not likely to help differentiate long-run effectiveness because all policies have similar medium-term outcomes regardless of the size of spillovers.

The clearest policy implication of our research is that to the extent that policymakers choose to subsidize innovative activity by firms, they should consider the full set of tax and regulatory policies that impact aggregate innovation through firm profitability. Taxing corporate profits or enacting regulations that make it more costly for firms to start up or operate has a significantly negative influence on innovation, undercutting the stimulative impact of R&D subsidization. The net effect may be to depress, rather than encourage, innovation by firms.

(intentional or inadvertent) of new innovations and knowledge created at just one location, firm or industry is central in justifying government subsidies for innovation. As a result, we want to know how important these spillovers are for the economy as a whole.

Economists who study these issues have generally followed two paths. The first is an effort to understand the impact of policy on innovation decisions taken *by individual firms*—the companies that develop new products and services or improve methods of production or service delivery. More relevant for overall economic well-being, the second looks at the *macroeconomy*, seeking to measure policy impact on a national level: To what extent do policies to encourage innovation generate broad economic growth? For the most part, these approaches have rarely intersected, leaving a significant gap in our knowledge of the mechanisms through which policy initiatives may or may not improve the economic well-being (or welfare) of Americans.

This policy paper reports on our effort to fill that gap by combining these two perspectives, the micro and the macro, thereby providing greater insight into several key considerations in innovation policy. (See "Aggregate Implications of Innovation Policy," Minneapolis Fed Staff Report 459, June 2011, online at minneapolisfed.org.) The overarching issue is: How do policies that affect firms' innovation costs and benefits impact aggregate output, productivity and welfare?

To answer this question, we have designed an economic model that is sufficiently detailed to capture the dynamic decisions of individual firms in response to innovation policy changes, yet is still mathematically manageable, allowing us to aggregate these many firm-level decisions and thereby gauge overall policy impact at the economywide (or "aggregate") level on output, productivity and economic well-being.

We break our analysis into two parts. In the first, we study which policies are most efficient in the long run in balancing their fiscal cost to taxpayers against their benefits in stimulating overall innovative activity by firms, including both those firms that are already operating when the policy is put in place and those that will enter under the new policy regime.

In the second part, we study how a policy-induced increase in innovative activity by firms impacts aggregate output, productivity and welfare (taking into account the fiscal cost to taxpayers of the policies used to stimulate that innovation) over both the long run and a medium-term horizon of 15 years.

Our research gives new answers to both of these questions.

Consider first the balance between the fiscal cost of various innovation policies and their effectiveness in stimulating innovative activities by firms. The standard analyses of fiscal efficiency of innovation policies attempt to fathom the many intricate details of the effects of the new policies on individual firms' decisions about changing investments.[3] By contrast, in our research, we embed a model of firms' innovation decisions in an overall model of the macroeconomy and show that with such a model, we can estimate the policy effectiveness of stimulating innovative activity simply by calculating the approximate impact of that policy on the profitability of new firms that might enter under this new policy.

Moreover, our results imply that, under some conditions, a broad set of innovation policies are all equally efficient: Two policies have the same impact on aggregate innovation, output and productivity in the long run if they have the same fiscal impact on taxpayers. These results provide a simpler procedure for evaluating the effectiveness of innovation policy and other government efforts to stimulate innovative activity and suggest that policies currently in place to stimulate innovation might also be dramatically simplified.

To gauge policy strength in encouraging innovative investments by firms, we focus on the size of policy impact on the potential profitability of new firms. By doing so, our research implies that innovation subsidies and tax preferences are only part of a much broader set of government policies with both positive and negative effects on firms' incentives to innovate. In particular, the negative impact of the corporate profits tax on incentives to innovate through its impact on potential profitability of new firms[4] may very well undo, at the aggregate level, much of the benefit of current direct federal support for R&D. In this respect, our research indicates that in terms of their effect on innovation investment, the current mix of federal subsidy and tax policies may negate with one policy the impact of others.

On the second question of how a policy-induced increase in innovative activity by firms affects long-run aggregate output, productivity and welfare (taking into account the transition from the status quo to the long run), our research indicates that it may be very difficult to reach definitive conclusions, given available data. Our model predicts on the one hand that a policy-induced change in innovative activity by firms may have a very large impact on output and productivity in the long run and on welfare, particularly if the spillovers from innovative activity are large. On the other hand, if spillovers are small, this may not be the case (and policies to stimulate innovative activity may not raise economic well-being).

Our research indicates that analysts may not be able to distinguish between those divergent long-run outcomes because our model's predictions for the macroeconomic response to an innovation policy over a reasonable time horizon, such as 15 years, look quite similar whether spillovers are very large or very small. Our model's simulations of the economy's medium-term response to a significant increase in innovation subsidies suggest that analysts working with real-world data would have difficulty obtaining reliable estimates of the magnitude of innovation spillovers for the economy as a whole and hence the implications of actual policy changes for welfare.[5]

Our research approach

To analyze the micro- and macroeconomic dynamics of innovation policy, we've built an economic model that is rich, yet tractable. By this, we mean that it combines the fundamental and detailed elements of innovation processes at a company level, but nonetheless allows us to generate estimates of the overall national economic impact of these firm-by-firm decisions as influenced by changes in government innovation policy. This policy paper describes our model and research in broad terms, shares the analytical and quantitative insights we've gained and then discusses implications of these findings for both research and policy.

Our research approach and this paper consist of several steps.

- First, we build a model that enables us, with a two-stage procedure, to assess the impact of changes in firm-directed innovation policies on macroeconomic output, productivity and economic well-being.
- We then use this procedure to establish several analytical results about the long-run response of a macroeconomy, through its microeconomic units, to innovation policy change.
- These results allow us, in a third step, to assess the relative and absolute size of the medium-run and long-run macroeconomic impact generated by several distinct real-world innovation policy options.
- These findings imply several directions for future policy, discussed briefly in a final section.

Our model

We use a dynamic general equilibrium model common to macroeconomic research. It includes households that work and consume according to their preferences and budget constraints, firms that invest and produce with specified technologies and objectives, and a government that has objectives, revenues and expenditures.

To this standard framework, we introduce a number of special features that allow us to analyze macroeconomic (also referred to as "aggregate") implications of innovation policy.

We build a model of monopolistically competitive[6] firms that engage in either *process* innovation, which will increase their productivity (a more efficient assembly line, as a mundane example), or *product* innovation, which enables them to create a new type of product (an iPad, if they're very lucky and skilled). More simply put, firms can invest in R&D to become more productive or expand the range of goods available to consumers.

To consider the impact of policies on the cost of innovative activity, our model includes a *research good* that firms use as an input for innovation. Research goods aren't unlimited. Producing them requires a combination of goods and labor (lab equipment and scientists) along with scientific knowledge that is freely available. Also, most crucially for purposes of our analysis, the production of the research good benefits from *innovation spillovers*—the knowledge and experience that researchers gain through innovation activities that neither they nor their company directly capture (at least financially). Through these spillovers, current innovative investments by firms have an added benefit to society of increasing the productivity of future R&D workers and thus reducing the cost to firms of future innovation.

In our model, we consider the i*mpact of a range of subsidies* financed by taxes collected from households and equaling aggregate fiscal expenditures by government. These subsidies—fairly abstract when we derive analytical results and later made concrete in our quantitative estimates—include a subsidy to variable profits from production, a subsidy to process innovation and a subsidy to product innovation. In addition, firms are taxed on their use of physical capital, essentially a negative subsidy.

While this brief description hardly does justice to a relatively complex model, it provides a sense of the key features that allow us to analyze the impact of innovation policy on both individual firms and the macroeconomy.

We discovered that a relatively straightforward procedure—a two-step algorithm, or sequence of formulas in which results from the first are inputs to the second—would provide approximate estimates of the long-run impact on macroeconomic outcomes of changes in innovation policy and thereby enable us to compare the relative and absolute magnitudes of the impact of various policy alternatives.

A two-step procedure

Detailed examination of the interaction of these features and the more standard variables in our model yields insight into what is (and isn't) fundamental to analysis of the macroeconomic impact of changes in innovation policy. We discovered that a relatively straightforward procedure—a two-step algorithm, or sequence of formulas in which results from the first are inputs to the second—would provide approximate estimates of the long-run impact on macroeconomic outcomes of changes in innovation policy and thereby enable us to compare the relative and absolute magnitudes of the impact of various policy alternatives.

The first step in this procedure is using a basic formula to measure the impact of policy changes on the profits an entrepreneur might expect from starting a new firm.

The second step is to then use the model's macroeconomic structure to infer long-run changes in aggregate output and wages that must result, in general equilibrium, to restore the incentives of entrepreneurs to create new firms or products in the face of the estimated change in expected firm profitability calculated in the first step.

In other words, the procedure gives us estimates of the new long-term level of macroeconomic outcomes that corresponds to whatever change in firm profits is generated by a new government innovation policy. And it does so without having to fathom the many intricate details of the new policy's effect on firms' decisions about changing investments, hiring, corporate structure and the like. We need only compute how the policy changes firm profitability—a far easier task.

This straightforward procedure (and the reasoning behind it) allows us to analyze more fully the implications of innovation policy changes. We do so in the next section, followed by an examination of the *quantitative* application of the procedure.

Analytical results

A central insight offered by our model and the algorithm just described is that a subsidy to all types of innovative activity has the same impact on macroeconomic outcomes as a direct subsidy to firm profitability. The reasoning is quite intuitive. Subsidizing a firm's innovative activities—in this case, by changing the price of the research good with a uniform subsidy to process and product innovation—lowers its costs, or equivalently, raises profits. Since profits here are the returns to innovation, supporting firm innovation through a subsidy has an identical impact on firm behavior and aggregates as a direct subsidy to firm profits.

We also find that, under some conditions, whether the subsidy is directed toward *process* or *product* innovation makes little difference in computing the effect on the *macroeconomy* as long as the impact on firm profitability is the same; this is because of dynamics that ensure that in macroeconomic equilibrium, with free entry of firms, companies will start up in an industry until doing so would no longer offer profits to entrepreneurs. (A policy directed specifically at either process or product innovation may have a dramatic impact on *firm*-level behavior, however, particularly on the innovative investments of existing firms.)

The zero-profit condition for entrepreneurs considering starting firms in a given industry limits the aggregate response of innovative investments by both existing firms and entrepreneurs contemplating a startup venture. This analytic insight is what is distinctive about our method for measuring the response of firms' innovative investments to a change in policy.

Previous research has often focused on the innovative response of existing firms only and neglected to consider that—in the long run, in general equilibrium—*the zero-profit condition for entrepreneurs creating new products is key* to assessing the overall response of the economy to the policy change. With this analytical insight, we argue that *regardless of how existing firm investments react to specific subsidies, the response of the macroeconomy will be the same.*[7]

In terms of policy, this implies further that, as we alluded to earlier, the details of firms' responses to changes in innovation policy are not of great importance for aggregate outcomes; beyond pure subsidization of profits, there is no special role for innovation policies. An example clarifies the implications of this argument. Consider the current design of the Research & Experimentation Tax Credit. This innovation policy sets out a complex set of rules by which a firm can gain a corporate tax credit for "qualifying research and experimentation expenditures" over and above a defined "baseline amount." The underlying idea is to reward existing firms only for new or incremental investments in innovation and to avoid subsidizing firms for innovation they would have done anyway.

Our research indicates that this policy focus on incremental expenditures at the firm level is misguided, since the impact on existing firms' investment is not the factor that determines the impact of the policy in the long run. Instead, it is the impact of the tax credit on the incentives of entrepreneurs to start new firms or introduce new products. Our results imply that the Research & Experimentation Tax Credit is, therefore, an administratively expensive way of offering a small reward to entrepreneurs who consider starting a new firm and spending money on R&D that qualifies for the credit somewhere down the line as their new firm grows. It would be more straightforward (and more efficient in terms of administrative costs) to subsidize firms in the relevant industry directly.

Quantitative results

These analytical results lead to the question of magnitude. How can we measure the effectiveness of various innovation policies in stimulating innovative investments given their fiscal cost to taxpayers? And how can we measure the impact of this induced innovation on aggregate productivity, output and welfare? We conducted two sorts of quantitative analyses. The first measured the *relative* impact of several innovation policy options. The second calculated the *absolute* size of the economic effect of parallel policy options.

Comparison of relative policy impact

To understand the effect of innovation policy on broad economic growth and welfare and to evaluate the relative efficacy of different policy options, a means of quantifying and comparing financial cause and effect—that is, cost and benefit—is essential. With our model, we show that, to a first-order (or ballpark) approximation, the relative impact of a policy change on firm profitability and on macroeconomic aggregates in the long run is proportional to the impact of the policy change on government fiscal expenditure.

In other words, to compare, roughly, how large an impact alternative innovation policy options will have both at the level of firm profits and on broad economic outcomes in the long run—GDP and productivity—we need only calculate how much that policy costs. The two figures aren't equal, just proportional, and the calculation is only a rough estimate, not a precise figure. But it means that to evaluate the relative merits of alternative policy options, we need only know their fiscal impact; the difficult task of gauging how millions of firms will respond to the policy isn't necessary.

To apply our results to actual policies in the United States, we looked at (1) the Research and Experimentation Tax Credit program, (2) federal spending on research and development and (3) the corporate profits tax. (Beyond the well-understood effects of the corporate profits tax on investments in physical capital, the tax influences innovation decisions in two ways: It affects variable after-tax profits generated from improved products or process, and firms may expense a portion of the cost of innovative activity and thus deduct these expenses from taxable profits. To the extent that firms are not able to fully deduct all of their expenses for innovation or are not able to carry forward all of the loss when attempts at innovation are unsuccessful, the net effect of the corporate profits tax is to reduce the profitability of starting a new firm or introducing a new product.)

Data from 2007 indicate that fiscal expenditure on the Research and Experimentation Tax Credit was \$10 billion. In the same year, federal spending on the five categories grouped into R&D by the Office of Management and Budget—basic research, applied research, development, R&D equipment and R&D facilities—totaled \$139 billion. (In contrast, business R&D spending in 2007 was far higher, about \$260 billion.) Comparing these two figures (and applying the appropriate discount factor since subsidies to product innovation are paid upfront while variable profits are received in the future), we can clearly see that the *long-term impact on aggregate output of federal R&D spending is far larger than the impact of the Research & Experimentation Tax Credit.*

Calculating the impact of the corporate profits tax—which raised \$445 billion in federal revenue in 2007—is more complicated because it depends on parameter values in a quantitative model that affect the physical capital-to-output ratio. But once parameters are chosen, we find that the long-run impact of the corporate profits tax (per dollar of revenue raised) exceeds that of innovation policies (per dollar spent) unless innovation spillovers are very large.

Hence, in our calibrated model, described below, reducing the corporate profits tax to collect \$100 billion less in revenue would have a comparable or even larger impact on innovation spending and aggregate output in the long run than increasing either the Research & Experimentation Tax Credit or federal R&D spending by \$100 billion, unless spillovers are very high. We thus conclude that *the corporate profits tax may very well be a relatively potent, counterproductive policy in terms of discouraging the long-run accumulation by firms of both physical, tangible capital and intangible capital (that is, patents, trademarks, intellectual property and the like).*

Comparison of absolute magnitude of policy impact

In a second quantitative exercise, we evaluated the absolute magnitude of both the long-run and medium-term impact on the macroeconomy of innova-

We found that none of the subsidies has significant impact on economic welfare if innovation spillovers are small. Output and productivity rise in the long run (and perhaps by a lot), but this increase comes at the cost of inefficiently high investments in innovation and low consumption by households in the transition from the present to the long run.

tion policies after putting some concrete figures into our model, giving it further realism by providing reasonable values for parameters such as the GDP growth rate, interest rate and capital depreciation. With this calibrated model, we measured the absolute magnitude of impact on GDP, welfare, productivity, research intensity and other economywide outcomes of two policies:

(1) A uniform subsidy to innovative activities (meaning that both process and product innovation would receive support).

(2) A subsidy to process innovation only.

In each case, the subsidy represented a fiscal expenditure of 3 percent of GDP, or about $420 billion in 2007 (similar to the revenue raised from corporate profits taxes that year). These are two typical policies aimed at stimulating innovation.[8]

Long-run response

In the long run, we find, innovation policies have an impact on the scale of firms' investments in innovation similar in magnitude to their fiscal impact, both relative to the level of GDP. Specifically, the *research intensity of the economy* (defined as the ratio of firms' spending on innovative activities to GDP) increases by roughly 3 percentage points of GDP in response to a subsidy of 3 percent of GDP. Moreover, this response of firms' innovative activity to innovation policy change is the same in the long run and roughly the same in the medium term regardless of the level of spillovers from innovative activity.[9]

Do these policies aimed at stimulating innovation increase consumers' welfare? The answer to this question is not obvious. At first glance, it appears that such a policy might not increase economic well-being—the taxes that consumers must pay to finance these innovation subsidies are roughly the same as the increase in firms' investments in innovation that result. In the absence of spillovers from firms' innovative activity, a policy of taxing households to pay for firms' investments in intangible capital is not likely to improve households' well-being. In the presence of spillovers, however, such a policy might bring substantial welfare benefits.

Our model confirms this logic. We found that none of the subsidies has significant impact on economic welfare *if* innovation spillovers are *small*. Output and productivity rise in the long run (and perhaps by a lot), but this increase comes at the cost of inefficiently high investments in innovation and low consumption by households in the transition from the present to the long run.

If spillovers are *large*, however, the subsidies have far greater impact on economic well-being. In fact, in this case, innovation subsidies of 3 percent of GDP can bring huge gains for households. The numbers from our model simulations below illustrate this point. We measure improvements in household economic welfare from policy changes by the amount that household consumption would have to be increased each and every year under the old policy to make households as happy as they would be with the consumption they attain under the new policy.

When we set our parameter for innovation spillovers at zero, the impact of the innovation subsidies on welfare is very close to zero—consumers would be just as happy with or without the innovation policy. In contrast, when we set our parameter for innovation spillovers close to its maximum possible value consistent with balanced growth, the impact of innovation policies on welfare is very large. Consumption under the old policy would need to rise by roughly 50 percent every year to attain the same level of household welfare as

Continued on page 43

Innovation from page 11

achieved in the equilibrium with innovation subsidies. Welfare gains like these are why Nobel Prize-winning economist Robert E. Lucas Jr. wrote that once one starts thinking about long-run growth and economic development, "it is hard to think about anything else" (p. 5).[10]

Our results on the long-run impact of innovation policies on aggregate output and productivity are also highly sensitive to our assumption for the parameter governing spillovers from innovative activity. When we set our parameter for innovation spillovers to zero, GDP is estimated to increase by a factor of only 1.03 (that is, by 3 percent) in the long run. In other words, in this case, the subsidies have little impact on either output or welfare.

But when we set the spillover parameter close to the maximum value consistent with balanced growth, the impact on GDP is much larger: It increases by a factor of 9.88 for policy 1 and 8.25 for policy 2. These nearly tenfold changes in GDP are comparable to the growth that the United States experienced from the beginning to the end of the 20th century and are brought about by a substantial, but perfectly feasible, level of innovation subsidies.

Clearly, *our model's implications for the long-run impact of a given change in policies vary tremendously depending on the assumed spillover parameter. If spillovers are large, there is a lot at stake for consumers in getting innovation policy right.*

Medium-term response

Our results on the impact of innovation policies on welfare and on output and productivity in the long run prompt the question: Can we use data on the response of the macroeconomy to changes in innovation policy over the medium term (say, 15 years) to figure out if spillovers from firms' innovative activities are small or large?

There is a large literature that attempts to answer this question, but, as Griliches (1988) and CBO (2005) discuss, the changes in the innovation intensity of the U.S. economy seen in the historical data are relatively small. It is therefore difficult to distinguish the effects on the macroeconomy of such small changes in R&D spending from the effects of all the other major factors at play—education, population growth and international trade, to name a few.

To shed light on the question of whether we might be able to measure economywide spillovers from innovative activity using available data even if we were to observe a large change in the innovation intensity of the economy arising from a change in innovation policy, we examined how the model performed over a shorter time frame, a 15-year medium-term period. The idea here is to understand transition dynamics—between now and the long-term equilibrium, *how* does the economy evolve, and *what factors are important* in that evolution? Again, and for all policies, we use the same subsidy size: a fiscal expenditure of about 3 percent of GDP.

Surprisingly, perhaps, we found that over this time frame, the two innovation policy options have a *similar impact* on economic growth *regardless of innovation spillover size*. In all cases, the cumulative factor increase in GDP in the 15th year is between 1.01 (or 1 percent) with no spillovers and 1.05 (5 percent) with high spillovers. Such small differences in GDP over a 15-year horizon would likely be difficult to discern in real-world data. Therefore, our results indicate that data on the response of GDP to innovation policy changes over the *medium* term will not shed much light on the size of such spillovers, suggesting that estimating policy outcomes over the *long* term will remain difficult, since an accurate measure of spillovers can't be obtained from shorter-term data.[11]

What explains the significance of spillovers for welfare?

The contrast in findings between long- and medium-term significance of innovation spillovers raises the question of why spillovers would have importance on innovation's macroeconomic impact *only* in the long run.

The intuition for this result is simply the idea of compound interest. Over the medium term, innovation policies have a similar impact on GDP growth regardless of the level of innovation spillovers. The real impact of spillovers comes only at longer time horizons. In the absence of spillovers, the boost to growth from innovation subsidies peters out relatively quickly and households are left paying roughly the same amount in taxes as the gain to innovation spending and the increase in GDP achieved. In contrast, if spillovers are large, the boost to the growth of GDP from increased investments in innovation lasts for a long time, well beyond the medium-term horizon, and innovation spillovers compound over time, bringing large benefits associated with a moderate boost to growth that lasts over 100 years.

Summary and implications for policy

We've established a benchmark model of innovation that provides a straightforward procedure for estimating relative magnitudes of long-run macroeconomic impact of a range of innovation policy options. The procedure gauges approximate policy impact on macroeconomic outcomes quite simply, through computing the government's fiscal expenditure on innovation policies.

The response of economic welfare and GDP over the long run to changes in innovation policy is highly sensitive to the size of innovation spillovers; welfare gains could vary between virtually no change and a 50 percent increase in equivalent consumption, depending on spillover size.

Unfortunately, we cannot accurately measure these long-run effects without accurate estimates as to the magnitude of innovation spillovers. Results from our model indicate, however, that even under ideal conditions, it should be very difficult to measure spillovers using data on medium-term response of the macroeconomy from changes in innovation policy. That is, evidence from the medium term is not likely to help differentiate long-run effectiveness because all policies have similar medium-term outcomes regardless of the size of spillovers.

What does this imply for policy?

The clearest implication of our research is that to the extent that policymakers choose to subsidize innovative activity by firms, they should consider the full set of tax and regulatory policies that impact aggregate innovation through firm profitability. Taxing corporate profits or enacting regulations that make it more costly for firms to start up or operate has a significantly negative influence on innovation, undercutting the stimulative impact of R&D subsidization. The net effect may be to depress, rather than encourage, innovation by firms. R

Endnotes

[1] This paper is based on: "Aggregate Implications of Innovation Policy," Minneapolis Fed Staff Report 459, June 2011. The authors thank Doug Clement for assistance in preparing this text.

[2] See "Strategy for American Innovation: Introduction." Also see Chairman Ben Bernanke's May 16, 2011, speech, "Promoting Research and Development: The Government's Role," for a discussion of the importance of innovation by firms to long-run growth and a summary of the questions regarding the rationale for, the effectiveness of and the impact of federal support for research and development that we address in this policy paper.

[3] Specifically, following the methodology developed by Hall and Jorgenson (1976) for physical capital, a standard approach is to first estimate the impact of a policy change on the "user cost of R&D" and then estimate the elasticity of firms' demand for R&D in response to such a policy-induced change in the user cost of R&D. See Hall and Van Reenen (2000) and CBO (2007) for examples of such analysis.

[4] See Gentry and Hubbard (2000) and Cullen and Gordon (2007) for a discussion of the mechanisms through which the U.S. tax structure reduces the incentives of entrepreneurs to start new firms.

[5] In this sense, our research casts doubt on the methods economists have previously used to measure the relationship between innovative activity by firms and aggregate productivity in the long run. See, for example, CBO (2005) and Hall, Mairesse and Mohnen (2009) for summaries of this research.

[6] A monopolistically competitive market combines characteristics of competition and monopoly. There are many buyers and many firms, with free exit and entry into industries, as under perfect competition. But consumers perceive sufficiently great nonprice differences (branding, for example) among similar products that producers can exercise a degree of control over pricing, as in a monopoly. Brand-name cereals and restaurants are textbook cases; laptop computers might be another example.

[7] Innovation policies in our model do impact the user cost of R&D and do have an impact on the innovative investments by incumbent firms that does depend on the responsiveness of these incumbent firms' innovative investments to changes in the user cost of R&D. This responsiveness, or elasticity, of R&D investments is not of first-order importance, however, in the calculation of how a change in inno-

vation policy affects the expected profitability of a new firm. For example, in calculating the impact on the expected profitability of a new firm from a change in a tax credit for R&D, what is of first-order importance is the change in taxes that a new firm can expect to pay given the investments in R&D that it had planned to undertake before the policy change was proposed. For small changes in policy, the additional accuracy gained by considering the impact on the expected profitability of new firms that arises from considering changes in policy and firms' investments simultaneously is necessarily very small.

[8] In the full paper, we also consider a third policy, a subsidy to physical capital to compare the impact of policies aimed at promoting firms' investment in intangible capital and those promoting investment in physical or tangible capital. We make this comparison to analyze the impact of the corporate profits tax, which is a combination of taxes on firms' profits from intangible and tangible capital.

[9] Our findings here are consistent with those summarized by Hall and Van Reenen (2000) on the effectiveness of fiscal incentives for R&D.

[10] See Lucas (1988).

[11] We note that substantial research has sought to establish a link between research intensity and output or productivity. This research has generally used regression analysis of disaggregated data at the firm or industry level. Unfortunately, this evidence is less than conclusive for answering the questions addressed by this paper.

First, many of these results are driven by long-term differences across firms or industries: Firms and industries that invest more in R&D also appear to have higher levels of productivity. It is not clear, however, how to interpret this observation. Klette and Kortum (2004), for example, argue that it should be accounted for by models with intrinsic factors that vary across firms and industries and that it does not necessarily indicate that a policy of stimulating further R&D would have a substantial impact on the aggregate economy.

Second, even in our model, a policy aimed at stimulating innovative investments by a *select group* of firms or industries can have a large impact in the short term on output and productivity that suggests spillovers are high even if *aggregate* spillovers are absent. It is a simple matter for a subset of firms or industries to invest in innovation and grow at the expense of the other firms or industries in the economy. Therefore, evidence of specific firm or industry responses to policy changes does not necessarily shed light on the central question of the *macroeconomic* response.

References

Bernanke, Ben S. 2011. "Promoting Research and Development: The Government's Role." Speech at the Conference on "New Building Blocks for Jobs and Economic Growth." Washington, D.C., May 16. Online at federalreserve.gov.

Congressional Budget Office (CBO). 2005. "R&D and Productivity Growth." Background Paper, June.

Congressional Budget Office (CBO). 2007. "Federal Support for Research and Development." A CBO Study, June.

Cullen, Julie Berry, and Roger H. Gordon. 2007. "Taxes and Entrepreneurial Risk-Taking: Theory and Evidence for the U.S." *Journal of Public Economics* 91(7/8): 1479-1505.

Gentry, William M., and R. Glenn Hubbard. 2000. "Tax Policy and Entrepreneurial Entry." *American Economic Review* 90 (2): 283-87.

Griliches, Zvi. 1988. "Productivity Puzzles and R&D: Another Nonexplanation." *Journal of Economic Perspectives* 2 (4): 9-21.

Hall, Bronwyn H., Jacques Mairesse and Pierre Mohnen. 2009. "Measuring the Returns to R&D." *National Bureau of Economic Research Working Paper* 15622.

Hall, Bronwyn H., and John Van Reenen. 2000. "How Effective Are Fiscal Incentives for R&D? A Review of the Evidence." *Research Policy* 29 (4/5): 449-69.

Hall, Robert E., and Dale W. Jorgenson. 1976. "Tax Policy and Investment Behavior." *American Economic Review* 57: 391-414.

Klette, Tor Jakob, and Samuel Kortum. 2004. "Innovating Firms and Aggregate Innovation." *Journal of Political Economy* 112 (5): 986-1018.

Lucas, Robert E. Jr. 1988. "On the Mechanics of Economic Development." *Journal of Monetary Economics* 22 (1): 3-42.