

NBER WORKING PAPER SERIES

LESSONS FROM PATENTS: USING PATENTS TO MEASURE
TECHNOLOGICAL CHANGE IN ENVIRONMENTAL MODELS

David Popp

Working Paper 9978

<http://www.nber.org/papers/w9978>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2003

The author thanks Robert Stavins and Ian Sue Wing for comments on an earlier draft. Remaining errors are the responsibility of the author. At various stages of this research, Brant Brown, Yixin Dai, Neelakshi Medhi, and Janna Matlack provided valuable research assistance. Financial support for the works cited within provided by DOE grant 593A3140217 and National Science Foundation grant SES-0001679. The views expressed herein are those of the authors and are not necessarily those of the National Bureau of Economic Research.

©2003 by David Popp. 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.

Lessons From Patents: Using Patents To Measure Technological Change in
Environmental Models
David Popp
NBER Working Paper No. 9978
September 2003
JEL No. Q33, Q30, Q40

ABSTRACT

When studying solutions to long-term environmental problems such as climate change, it is important to consider the role that technological change may play. Nonetheless, to date few economic models of climate change explicitly model the link between policy and technological change. There is a growing body of evidence that the incentives offered by prices and environmental regulations have a strong influence on both the creation and adoption of new technologies. In several recent papers, I have used patent data to examine the links between environmental policy and technological change. In addition, I have used the results of this research to calibrate the ENTICE model (for ENdogenous Technological change) of climate change, which links energy-related R&D to changes in the price of carbon. Drawing on my experiences from empirical studies on innovation and from modeling the climate change problem, in this paper I review some of the key lessons from recent empirical work using patents to study environmental innovation and diffusion, and discuss its implications for modeling climate change policy. I conclude by offering suggestions for future research.

David Popp
Department of Public Administration
The Maxwell School
Syracuse University
400 Eggers Hall
Syracuse, NY 13244-1090
and NBER
dcpopp@maxwell.syr.edu

Debates about the sustainability of environmental policy (or the lack thereof) often revolve around the potential effects that technological change may have. Nonetheless, most economic models of long-term environmental policies such as climate change treat technology as exogenous. Since these policies are likely to have a large impact on the pace and direction of technological change, such models miss the important link between policy and innovation. One reason for this omission is that, until recently, few empirical studies of innovation and environmental policy existed.¹

My own research has focused on filling this gap. In several recent papers (Popp 2001, 2002, 2003a), I use patent data to examine the links between environmental policy and technological change. Patent data offer several advantages when studying environmentally-friendly technological change. Detailed patent classification systems make identifying advances in narrowly-defined technological fields straightforward. Furthermore, since inventors can apply for patents in multiple countries, patents can also be used to track the diffusion of technologies across countries. I have used the results of this research to calibrate the ENTICE model (for ENdogenous Technological change) of climate change. The ENTICE model is a variant of the well-known DICE model of climate change (Nordhaus 1994) that links energy-related R&D to changes in the price of carbon. Drawing on my experiences from empirical studies on innovation and from modeling the climate change problem, in this paper I review some of the key lessons from recent empirical work using patents to study environmental innovation. Using the ENTICE model as a an application of these lessons, I discuss the implications for modeling long-run environmental problems, and how ignoring the lessons learned from empirical work helps explain differences found in various papers. I conclude by offering suggestions for future research.

I. Theoretical Background

The process of technological change proceeds in stages (Schumpeter 1942). First, an idea must be born. This stage is known as *invention*. New ideas are then developed into commercially viable

¹ Although recent models, such as those presented in this special issue, help to fill this gap in the research, the wide range of results presented in these models suggest there is still much work to be done.

products. This stage is *innovation*. Typically, these two stages of technological change are studied together. Finally, to have an effect on the economy, individuals must choose to make use of the new innovation. This adoption process is known as *diffusion*. At each stage, incentives, either in the form of prices or regulations, will affect the development and adoption of new technologies. This section reviews the theoretical literature linking economic incentives to the process of technological change.

A. Induced Innovation

The idea that the creation of ideas through invention and innovation will be influenced by economic incentives is not new. The concept of *induced innovation* was first introduced by John Hicks (1932), who noted that changes in factor prices would lead to innovation to economize on usage of the more expensive factor. Theoretical papers by Ahmad (1966), Kamien and Schwartz (1968), and Binswanger (1974, 1978a, 1978b) formalized Hicks' notion and established the theory of induced innovation. Binswanger notes that not only will an increase in the cost of a factor of production lead to an increase in R&D focused on reducing use of that factor, but also that changes in the productivity of such R&D will affect inducement. This suggests the importance of considering the marginal productivity of research. As I will discuss later, there is evidence that, within a given field, the productivity of research declines over time. As a result, an increase in energy prices should lead to an immediate increase in R&D on energy conservation, but this increase need not be long-lived.

In the environmental literature, the relationship between innovation and policy has been explored under two broad themes.² Papers such as Magat (1978), Milliman and Prince (1989), and Fisher *et al.* (1998) use theoretical models to compare the effects of various environmental policy mechanisms (e.g. command and control regulation, emissions taxes, or tradable permits) on environmentally-friendly innovation. In general, these papers predict that market-based policies, such as a tax or tradable permit, will induce more environmentally-friendly innovation than a command and control policy.

² A comprehensive survey of the literature on technological change and the environment may be found in Jaffe, Newell, and Stavins (2003). Jaffe, Newell, and Stavins (1999) discuss the implications of energy-efficiency technologies on climate policy.

Although economists readily recognized the potential importance of linking technological change to climate policy, modeling such links was made difficult by the lack of empirical studies on the magnitude of such induced innovation effects. Recently, as measures of innovative activity such as patents have become more readily available, empirical economists have begun to estimate the effects that prices and environmental policies have on environmentally-friendly innovation. Using international patent data, Lanjouw and Mody (1996) find a correlation across nations between patenting in environmental fields and pollution abatement expenditures. Jaffe and Palmer (1997) look at the same relationship across US industries. Comparing pollution abatement data and innovation data across these industries, they estimate a positive elasticity of pollution control R&D with respect to pollution control expenditures of 0.15. However, they do not find evidence of a relationship between patenting and pollution abatement expenditures.

In each of these studies, pollution abatement expenditures serve as a proxy for the stringency of environmental regulation. Other papers use energy prices and related regulations as the mechanism that induces innovation. Although the observed price changes might not be policy-related, the results can also be applied to situations where policy affects prices, such as a carbon tax. Newell *et al.* (1999) use an approach closely related to hedonic techniques to study the effect of both energy prices and energy efficiency regulations on technological advances in energy efficiency for air conditioners and natural gas water heaters. They find that energy prices have the largest inducement effect, although they do not estimate a price elasticity. In a series of studies (Popp 2001, 2002), I use patent data to estimate both the relationship between energy prices and innovative activity and the effect that this induced knowledge has on energy consumption. The results of these papers are discussed in greater detail in the subsequent sections.

B. Diffusion

Of course, induced invention and innovation is only half the story. Technological advances are of little use unless society makes use of the innovation. Thus, diffusion is also important. Diffusion is likely

to play a particularly important role for problems dealing with long-term consequences, such as climate change. Most innovation takes place in highly industrialized countries. In 1998, 85 percent of all R&D done in the OECD countries was done in just 7 countries. In fact, roughly 44 percent of all OECD R&D was done in the United States alone. (National Science Board, 2002) For example, as policymakers pay increased attention to potential increases in fossil fuel consumption for fast-growing countries such as China, it is important to not only pay attention to how policy will help induce the development of new technologies, but how policy can help encourage the adoption of these technologies in developing countries.

Studies of the diffusion of individual technologies consistently find that diffusion is a gradual process. Typically, the rate of diffusion can be represented by a sigmoid, or “S-shaped” curve over time: the rate of adoption rises slowly at first, speeds up, and then levels off as market saturation approaches. Figure 1 illustrates such a curve, with market penetration plotted on the y-axis, and the passage of time plotted on the x-axis.

Traditionally, researchers have used one of two approaches to modeling diffusion of a new technology. The *epidemic* model of diffusion proposes that information is the primary factor limiting diffusion. Adoption is slow at first, as few people (or firms) know about the technology. However, as more people adopt the technology, knowledge of the technology spreads quickly, leading to a period of rapid adoption. Economists often use the analogy of a contagious disease to describe this period of adoption – the more people “infected” by the technology, the more likely that others will also be “infected”. Eventually, few potential adopters remain, as nearly everyone has adopted the technology, so that the rate of adoption levels off again. Using this framework, Griliches (1957) noted that the rate of diffusion is at least partially determined by economic factors, such as the expected rate of return for adoption. Other work using the epidemic model, such as Mansfield (1968), Davies (1979), and Oster (1982), typically focus on firm characteristics, such as firm size, to explain variations in the rate of diffusion. In this framework, environmental policy enters as one of the economic factors affecting adoption.

The second approach to studying diffusion focuses is the *probit* model (David 1969). The probit model focuses on heterogeneity among firms. In this model, firm heterogeneity leads to a distribution of expected return from adopting the new technology. Since adoption is costly, only firms above a threshold great enough to justify the costs of adoption will choose to adopt the technology at any given time. Over time, the technology gets cheaper, and its quality improves, so that more firms cross the adoption threshold.³ In this framework, environmental policy enters by increasing the expected returns from adopting environmentally-friendly technology, and thus increasing the likelihood that firms will cross the threshold.

Although there is a broad economics literature on diffusion, both at the national and international level, few articles link diffusion and environmental policy. Rose and Joskow (1990) find that the adoption of fuel-saving technology by U.S. electric utilities is positively correlated with fuel prices. Jaffe and Stavins (1995) find that higher energy prices lead to greater use of insulation in new home construction, but that the costs of installation are a more important consideration. Both Gray and Shadbegian (1998) and Kerr and Newell (2001) show that environmental regulations increase the probability of adopting environmentally-friendly technologies, and that the response varies by firm characteristics.

Similarly, at the international level, there is a large body of research on technological diffusion in general (surveyed by Keller 2001), but little work focusing specifically on environmental technologies. One exception is Lanjouw and Mody (1996). Using patent data from the US, Japan, Germany, and 14 low-and middle-income countries, they find that environmentally-friendly innovation increases as pollution abatement cost expenditures in the country increase. For the US, Japan, and Germany, the majority of these patents are typically domestic patents. For the developing countries, the majority of these patents come from foreign countries, highlighting the importance of diffusion. Given the importance

³ Note that the two models suggest different interpretations of the gradual rate of diffusion. The epidemic model suggests that adoption generates positive externalities, and is thus slower than would be optimal. Conversely, the probit model suggests that gradual diffusion is optimal, as differences in adoption decisions simply result from profit-maximizing decisions by heterogeneous firms.

of diffusion for bringing new environmental technologies to developing countries, understanding the general lessons from empirical work on international diffusion is important. These lessons are discussed in Section V.

II. Using Patent Data to Study Environmental Technological Change

Several possible data sources exist for studying technological change and the environment. R&D data offer a straightforward measure of innovative activity. Diffusion studies often make use of adoption data for specific technologies. In comparison, patent data are not as straightforward. Patent counts should be expected to increase as R&D activity increases, but the correlation need not be exact. Furthermore, the existence of a patent does not mean that the technology has been adopted. Indeed, studies of the economic value of patents find that most patents have little commercial value, suggesting that adoption of most patented inventions is not widespread (see, for example, Lanjouw *et al.* 1998). Nonetheless, patent data offer several advantages when studying technological change and its effect on the environment.

Unlike more aggregate data such as R&D expenditures, patents provide a detailed record of each invention. From the bibliographic data on a patent, the researcher can learn the identity and home country of the inventor, read a description of the invention, and see references to earlier patents. Using this data, it is possible for researchers to collect data in highly disaggregated forms. Whereas R&D data are typically available only for specific industries or general applications,⁴ patent classifications can be used to distinguish between different types of R&D at great detail, such as air pollution control devices designed to reduce NO_x emissions versus devices designed to control SO₂ emissions.⁵ In addition, economists have found that patents, sorted by their date of application, provide a good indicator of R&D

⁴ For example, in the US, R&D data is available from 1972-1994 for air pollution control, but it is not broken down by pollutant.

⁵ For example, US patent classes 423/235-423/239 pertain to control of “nitrogen or (a) nitrogenous component”, and patent classes 423/242 – 423/244 and 423/569 – 423/570 pertain to control of sulfur compounds. Using patent databases, it is possible to download all patents in these classes.

activity (see, for example, Griliches 1990). As a result, patent counts not only serve as a measure of innovative output, but are indicative of the level of innovative activity itself.

Furthermore, patent data are available from many different countries. Sources such as the European Patent Office make available data from virtually every country.⁶ Of course, such data can be used to examine levels of innovative activity across countries. In addition, unlike R&D data, international patent data can be used to track patterns of diffusion. The legal protections granted by a patent only apply in the country in which the patent has been granted. Thus, inventors who intend to market a product in several countries may choose to patent in each potential market. Economists have used patent families – sets of patents related to the same invention that have been filed in various countries – to track diffusion of knowledge across countries (e.g. Lanjouw and Mody 1996).

Finally, when a patent is granted, it contains citations to earlier patents that are related to the current invention. The citations are placed in the patent after consultations among the applicant, his or her patent attorney, and the patent examiner. It is the applicant's responsibility to list any related previous patents of which he or she is aware. In addition, the examiner, who specializes in just a few patent classifications, will add other patents to the citations, as well as subtracting any irrelevant patents cited by the inventor. Patent citations narrow the reach of the new patent by placing the patents cited outside the realm of the current patent, so it is important that all relevant patents be included in the citations.⁷ For the same reason, inventors have an incentive to make sure that no unnecessary patents are cited. As a result, the previous patents cited by a new patent should be a good indicator of previous knowledge that was utilized by the inventor.⁸

⁶ Details about the data available through the esp@cnet, the European Patent Office's interface for world patent data, can be found at <http://ep.espacenet.com/>.

⁷ "Outside the realm" means that the patent holder cannot file an infringement suit against someone whose invention infringes on qualities of the patented invention that were also included in patents cited by the patent holder.

⁸ The key assumption here is that a citation made to a previous patent indicates a flow of knowledge from the cited patent to the citing patent, so that patents cited more frequently are considered more valuable to future inventors. Jaffe, Fogarty, and Banks (1998) investigate the validity of this assumption, using evidence from citations made to NASA patents. They conclude that, although there is noise in the citation process, aggregate citation patterns represent knowledge spillovers, although the spillover may be indirect.

In recent years, several economists have used patent citations as a measure of knowledge flows. Papers using citation data to capture the flows of knowledge across inventors (such as across institutions or across nations) include Jaffe, Fogarty, and Banks (1998), Jaffe and Trajtenberg (1996), Caballero and Jaffe (1993), Jaffe, Henderson, and Trajtenberg (1993), and Johnson and Popp (2003). Similarly, Popp (2002) uses patent citation data to proxy for knowledge flows across time, by using citation-weighted patent counts as a measure of knowledge available for future inventors to build upon.

Nonetheless, when working with patent data, it is important to be aware of its limitations. The existing literature on the benefits and drawbacks of using patent data is quite large.⁹ Most importantly, the quality of individual patents varies widely. Some inventions are extremely valuable, whereas others are of almost no commercial value. This is partly a result of the random nature of the inventive process. Accordingly, the results of studies using patent data are best interpreted as the effect of an “average” patent, rather than any specific invention.

In addition, although the decision to file a patent obviously follows from the decision to perform R&D, not all successful research results are patented. In return for receiving the monopoly rights inferred by a patent, the inventor is required to publicly disclose the invention. Rather than make this disclosure, inventors may prefer to keep an invention secret. Surveys of inventors indicate that the rate at which new innovations are patented varies across industry (Levin *et al.* 1987). As such, this propensity must be controlled for.¹⁰

⁹ Griliches (1990) provides a useful survey.

¹⁰ For example, in Popp (2002), which studies advances in energy technologies, I control for changes in the propensity to patent by using the percentage of all successful domestic patent applications per year in each technology field as the measure of innovative activity. Policy changes that affect patenting in *all* patent classifications would lead to a change in both the total number of patent applications and the number of pollution control patent applications in a given year. For example, a change in patent law that increases the propensity to patent by five percent would increase both the numerator and the denominator of my measure of innovation by five percent. Although this still requires that the propensity to patent did not vary independently *within the pollution control industry*, such an assumption is not nearly as strong as assuming the overall propensity to patent has remained constant.

III. Key Lessons from Empirical Studies of Innovation

Much of my research has focused on using patent data to study the links between policy and environmentally-friendly innovation. In my empirical work, I aim both to better understand how policy affects the evolution of technology, as well as to develop estimates of the magnitudes of these effects. I use these estimates to calibrate a simulation of climate change policy with induced technological change. In this section I offer several key lessons from this research. Section IV follows with the implications of this work for climate policy.

Lesson #1: Innovation responds quickly to incentives

Perhaps the most striking finding in my empirical work is the speed at which innovative activity responds to incentives. Consider, for example, Figure 2, taken from Popp (2002), which shows how patenting activity for various energy technologies responded to changes in energy prices. As is typical in the economics of innovation literature, the patents in this figure are sorted by their application date. As noted earlier, patents sorted by application date are a good indicator of R&D activity. Technologies such as solar energy experienced large jumps in patenting activity immediately following the first energy crisis. For example, in 1972 there were just 10 solar energy patents. This figure jumps to 36 in 1973, 104 in 1974, and 218 by 1975.

Using a distributed lag model, Popp (2002) estimates the elasticity of energy patenting activity with respect to energy prices for 11 different technologies. The distributed lag model is consistent with an adaptive expectations model of prices, in which expected future prices depend on a weighted average of past prices, so that

$$(1) \quad P_{E,t}^* = P_{E,t} + \lambda P_{E,t-1} + \lambda^2 P_{E,t-2} \cdots + \lambda^{t-1} P_{E,1},$$

where $P_{E,t}^*$ is the expected price at time t , and $P_{E,t}$ represents actual energy prices in year t . The adjustment coefficient, λ , represents the weights placed on past observations. Popp (2002) finds a long-run energy R&D elasticity of 0.354, and an adjustment coefficient of 0.829. As such, the mean lag occurs

in 3.71 years, and the median lag in 4.86 years. The interpretation here is that over one-half of the full effect of an energy price increase on R&D will have been experienced after just five years. Thus, prices (or other regulations that increase the cost of using fossil fuels) can be expected to stimulate new research quickly.

Lesson #2: Innovation in a given field experiences diminishing returns over time

Referring again to Figure 2, note that not only does energy R&D respond quickly to energy prices, but it also drops off more quickly than energy prices. Energy prices do not reach a peak until 1981. Nonetheless, patenting activity in most technologies peaks in the late 1970s. Had the returns to energy R&D remained constant over time, we would expect patenting activities in these fields to remain high until prices began to fall. These results suggest that the possibility of diminishing returns to research should be explored.

Popp (2002) shows that the returns to R&D are an important determinant of the level of innovative activity. Inventors “stand on the shoulders” of their predecessors. As a result, the quality of the existing knowledge on which an inventor can build is an important, positive contributor to the level of innovative activity in a given year. I use patent citations as evidence of the existing state of technology when an invention was completed. Citations to an earlier patent suggest that the previous patent provided technological knowledge upon which the current inventor could build. Frequent citations to a patent provide evidence that the knowledge embodied in that invention has been particularly useful to other inventors. To verify the importance of the existing knowledge stock on innovative activity, I use citation data to create stocks of patented knowledge, where patents in the stock are weighted by their propensity to be cited. I find that the stocks have a significant positive effect on energy R&D activity.

Furthermore, patterns of citation suggest that diminishing returns are important. I find the likelihood of citations to new energy patents falls over time, suggesting that the quality of knowledge

available for inventors to build upon also falls.¹¹ The intuition here is that, as more and more discoveries are made, it gets harder to develop a new innovation that improves upon the existing technology. Since the quality of the knowledge stock is an important determinant of the level of innovative activity, decreasing quality of the knowledge stock over time means that diminishing returns to R&D investment will result in lower levels of induced R&D over time.

In addition, although diminishing returns are important, such returns should not vary monotonically. Rather, as it is prior research that affects the potential success of future inventors, the returns to research should vary along with the quality of the existing pool of research. To verify the value of using patent citation data to measure the returns to research, Popp (2002) also includes regressions in which the stock of knowledge is replaced by a time trend. If diminishing returns proceed monotonically over time, a negative time trend should work as well as the weighted knowledge stocks. That, however, is not the case. These regressions prove unreliable. In fact, the elasticity of energy R&D to energy prices appears negative when a time trend is used in place of the knowledge stocks. Since diminishing returns are a bigger problem when the level of energy R&D is highest, not controlling for them counteracts the positive effect of prices on energy R&D.

Finally, it is important to note that the notion of diminishing returns within a given field is not inconsistent with the general notion that there are increasing returns to research. Here, diminishing returns refers to the expected returns on the *inputs* of the R&D process – the likelihood of success from additional R&D dollars spent in the energy field falls as more dollars are spent. When economists refer to increasing returns to research, they are referring to the returns on the *output* of research. These increasing returns result from the public goods nature of knowledge – once an invention is created, it can benefit everyone, and thus any individual inventor is likely to capture only a subset of the social value of his or her invention. Furthermore, as the expected returns to research in any one field decrease, we would expect researchers to shift more resources to more promising fields. Thus, diminishing returns over time

¹¹ Note that since the probability of a patent being cited depends not only on the quality of the patent, but also on the number of patents that follow, it is important to look at probability of citation, rather than raw citation counts.

within a field need not result in diminishing returns for all research over time. However, when focusing on innovation in a single field, such as energy-saving research in a climate change model, including diminishing returns is important.

Lesson #3 – A Time Trend is Not a Substitute for Technological Change

Typically, papers that attempt to assess the impact of technological change have done so using a time trend to control for technological advances. Unfortunately, such results may be misleading when translated into climate models incorporating induced innovation. The very notion of induced innovation suggests that changes such as new policies or price shocks should change the nature of R&D. Using a time trend to proxy for technological change does not allow for such a reaction. In addition, using a time trend only captures the *overall* effect of technological change. For example, low energy prices make it likely that energy use per unit output will increase, making technological change appear energy-using. That does not mean, however, that no energy saving technologies are created and adopted during periods of low energy prices.

Examples of the problem of interpreting time trends can be found by comparing empirical studies from different eras. In a study of U.S. industrial energy consumption from 1958 to 1974, Jorgenson and Fraumeni (1981) found that technological change was energy-using – energy use per unit output increased over time. Of course, the time period of their data would not include any of the energy saving innovations developed after the energy crises of the 1970s. Indeed, more recent work using a time trend to capture technological change finds that technology is energy saving. Examples include Berndt *et al.* (1993), Mountain *et al.* (1989) and Sterner (1990). Similarly, two empirical studies on the effect of new flue gas desulfurization (FGD) technologies use time trends to reach different conclusions on technological change. Looking at scrubbers installed by 1992, Bellas (1998) finds no significant evidence of technological change in abatement technology. In contrast, Carlson *et al.* (2000) find that about 20%, or \$50, of the change in marginal abatement costs that have occurred from 1985 to 1995 can be attributed to technological change. Such a change would be predicted by economic theory predicting market based

policies induce more technological change, as Bellas' data includes few scrubbers installed after passage of the 1990 CAA, and none installed after trading in SO₂ permits began. As shown in Popp (2003a), one reason for this change is that the nature of R&D shifted after permit trading began.

To illustrate the use of patents to estimate the effect of technological change, Popp (2001) uses energy patents to estimate the effect of new technology on energy consumption. In that paper, I use patent data to create stocks of knowledge. First, I match energy patents with the industries that use the inventions. Then, I create stocks of energy knowledge, which are used as an explanatory variable in a system of cost functions for 13 energy intensive industries. The knowledge stocks are defined as:

$$(2) \quad K_{i,t} = \sum_{s=0}^{\infty} e^{-\beta_1(s)} (1 - e^{-\beta_2(s+1)}) PAT_{i,s},$$

where $K_{i,t}$ represents the knowledge stock in industry i at time t , s represents an index of years up to and including year t , $PAT_{i,s}$ represents the number of energy patents used by industry i in year s , β_1 represents a rate of decay, and β_2 is a rate of diffusion. The stocks allow me to estimate the effect of a new patent on energy consumption. By combining these results with the induced innovation results discussed above, I calculate the effect of induced innovation as the combined effect of all new patents induced by a price change. I find that induced innovation accounts for approximately one-third of the change in energy consumption that occurs after a price change. Furthermore, the estimated elasticities of energy use with respect to price found in that paper are lower than typically found, as they include only the effect of factor substitution. By comparison, when the regressions from Popp (2001) are run using only a time trend to represent technological change, the resulting energy price elasticities are consistent with those found in other studies, and are comparable to the combined induced innovation elasticity and energy price elasticity estimated using knowledge stocks. Table 1, taken from Popp (2001), illustrates these results for selected technologies.

The table shows the short-run breakdown in the effects on energy consumption from a change in energy prices. The elasticity of energy use with respect to induced innovation is the percent change in energy consumption resulting from the new technologies induced by a one percent change in energy

prices. The elasticity of energy use with respect to price is the change in energy consumption resulting from factor substitution. Column three shows the total effect of a change in energy prices. It is the sum of the elasticities in columns one and two. Column 4 presents the percentage of the total elasticity that is due to induced innovation. Finally, column 5 gives the elasticity of energy with respect to price from a regression that excludes the knowledge stocks, so that only a time trend is used to capture technological change. Note that these elasticities are comparable to the combined elasticities presented in column 3. One implication here is that even models calibrated on typical estimates of energy price elasticity include some effects of technological change in them. Incorporating endogenous technological change into these models reveals the *marginal* effect of additional innovation induced by policy, but not the total effect of all technological innovations.

Lesson #4 – The Social Returns to Environmental Research are High

This lesson is not unique to environmental research. Consistently, economists studying the returns to research have found that the social returns are higher than the private returns to R&D. The reason for this is that the knowledge created by such research is a public good. Once new knowledge is publicly available, it can be applied repeatedly without decay. Moreover, without public policy, such as intellectual property rights, others cannot be excluded from making use of the new knowledge. As such, firms are unable to capture the entire social value of their research investment, and thus underinvest in research. Examples of such studies include Mansfield (1977, 1996), Pakes (1985), Jaffe (1986), and Hall (1995).

My research of environmental and energy technologies confirms that the social returns to research in these fields are comparable to that in other fields. As such, it provides a useful guide for calibrating the returns to R&D in climate policy models. In Popp (2001), I examine the effect of new knowledge, measure by patents, on energy consumption in 13 energy-intensive industries. The median patent leads to \$14.5 million dollars in long-run energy savings. In comparison, these industries spend an average of \$2.25 million of R&D per patent. Using similar techniques, Popp (2003a) examines the effect

of flue gas desulfurization (FGD) unit innovation on coal-fired electric power plants. Included in that paper is an estimate of the cost-savings that result from innovation. A single patent provides a present value of \$6 million in total cost savings across the industry. On average, about \$1.5 million of R&D are spent per patent granted.¹²

Lesson #5 – The type of policy used affects the nature of new innovations

Theoretical papers concerning the innovation incentives offered by various environmental policies consistently find that market mechanisms, such as emissions fees or tradable permits, offer greater incentives for innovation. Although most empirical work verifies these predictions, it is important to note that the superiority of market mechanisms for encouraging innovation depend not only on the *level* of innovation induced, but also on the *quality* of innovation induced.

As an example of this, Popp (2003a) finds that the level of innovation for FGD units was actually higher before tradable sulfur dioxide (SO₂) permits were introduced by the 1990 Clean Air Act (CAA). However, the nature of innovation changed after passage of the Act. Before the 1990 Clean Air Act, most new coal-fired electric utilities were required to install FGD units with a removal efficiency of 90%.¹³ Since installation of FGD units was mandatory, innovation focused on reducing the operating costs of these units. However, since there were no incentives for firms to exceed the 90% limit, innovation had no effect on the removal efficiency of FGD units.

In contrast, by offering fewer permits, the 1990 CAA required greater SO₂ emissions reductions, but did not specify how those reductions be met. While some firms chose to use more efficient FGD

¹² Two caveats are important. First, these studies focused on energy-intensive industries, where the likelihood of significant savings was high. Care should be taken in generalizing the results to other, less energy-intensive sectors. At the macro level, this implies that the rate of return may be lower if significant energy R&D efforts are directed at less promising research areas. Furthermore, this lesson should not be read as implying that public R&D dollars should be non-discriminately directed towards energy research, because past successes do not guarantee future high rates of return. Private rates of return are high because firms tend to direct research towards more promising projects.

¹³ Specifically, plants needed to either achieve emission rates less than 1.2 lbs SO₂/million Btu heat input *and* reduce 90 percent of their potential SO₂ emissions (a 90% removal efficiency) or achieve emission rates less than 0.6 lbs SO₂/million Btu heat input *and* achieve 70% removal efficiency. Potential SO₂ emissions are calculated based on the sulfur content of the fuel burned.

units, others were able to achieve reductions using alternative techniques, such as low-sulfur coal (Ellerman *et al.* 1997). While the overall level of FGD innovation was lower after passage of the 1990 CAA, the post-1990 innovations had a more positive environmental effect. During the 1990 CAA regime, the knowledge created each year increased the removal efficiency of FGD units by 1.71%. Indeed, in 1995 (the first year of trading in the SO₂ permit program), 11 new FGD units with a removal efficiency of 95% or higher were installed.

One possibility for the lower level of innovation is that fewer firms need to innovate in the post-90 regime. Regulations requiring a specific technology gave every affected firm incentive to innovate in ways that lowered the costs of compliance. However, the greater flexibility of permit trading offers firms several ways to comply. Some firms could lower pollution by switching to cleaner coal. Other firms, if necessary, could simply purchase permits from firms with lower marginal abatement costs. Only a subset of firms found that installing more efficient scrubbers to be the best method for meeting new pollution reduction targets.

Nonetheless, because of the public goods nature of new knowledge, having fewer innovating firms, but more promising innovations developed, is more beneficial to society. Once the know-how to produce 95% efficient scrubbers exists, it can be adopted by other firms.¹⁴ As an analogy, consider research to develop more fuel efficient cars. Society would benefit more from one firm that develops the ability to reduce emissions completely by developing a cost-effective solar-powered car, as opposed to several firms developing cars with slightly greater fuel efficiency. Given this, it is important for policymakers to be cognizant of not only the level of research induced by policy, but how different policies may shape the *outcome* of this research.

¹⁴ This assumes that there is a regulatory incentive to do so. Compared to emissions fees, this is a weakness of tradable permits, as the total emissions level remains fixed even after a more promising technology is found, as adopters of the new technology are able to sell unused permits to other firms.

IV. Implications for Climate Change

Despite the growing evidence that environmental policy influences the direction of technological change, few climate change models directly incorporate links between policy and technology.¹⁵ Those that do typically model innovation in one of two ways. *Bottom-up* models include a detailed specification of energy systems. These models typically do not include detailed modeling of the overall macroeconomy, and typically model induced technological change in a learning-by-doing framework, in which the costs of various technologies decrease with experience. Examples include Gerlagh and van der Zwaan (2003), Manne and Richels (2002), Grubler and Messner (1998), and Messner (1997). *Top-down* models focus on the links between environmental policy and macroeconomic performance. Endogenous technological change in these models typically comes through accumulated investment in research and development (R&D). Recent models of this nature include Goulder and Schneider (1999), Nordhaus (2002), and Buonanno *et al.* (2003).

However, none of the existing models make use of empirical estimates on the nature of technological change to calibrate the model. In recent work, I have begun to use the results of these empirical studies to explore the effect of induced technological change (ITC) in models of climate change. I use the estimates obtained in previous work to guide the development and calibration of a top-down model of endogenous technological change that is supported by empirical evidence. My work builds on the well-known DICE model of climate change (1994). The DICE model is a dynamic growth model of the global economy that includes links between economic activity, carbon emissions, and the climate. Because it includes both costs and benefits of climate change, it allows the modeler to simulate optimal paths for control of carbon emissions. I modify the basic DICE model to include a fossil fuel

¹⁵ The results of Chakravorty, Roumasset, and Tse (1997) show why considering the link between policy and technological change is important. They present a simulation showing that the development of alternative fuels lowers the projected temperature increases caused by global warming from 3-6 degrees centigrade to just 1-2 degrees centigrade. Unfortunately, they do not model the development of such technologies, but rather assume that they continue to develop at historical rates. Since much of the development of alternatives such as solar energy occurred during the period of high energy prices during the late 1970s, assuming that historical rate of technological progress will continue without policy stimulation is unrealistic, and may lead to overly optimistic results.

sector and an energy research sector.¹⁶ The modified model is called ENTICE, for ENdogenous Technological change.

Although the DICE model is a greatly simplified macroeconomic representation of the global economy and the environment, the relative transparency of the model makes sensitivity analysis straightforward. Existing models of technological change and the environment include differing assumptions about features such as imperfect research markets, the potential crowding out effects of energy R&D, and key elasticities. The ENTICE framework allows detailed sensitivity analysis of these assumptions, helping to rectify the differences found in various models. To illustrate the importance of the lessons presented earlier, in this section I highlight key results from this still ongoing work.

Implication #1: The Long-Run Effects of Technological Change are Important

To compare welfare with and without induced technological change, I calculate the net economic impact for an optimal climate change policy with and without ITC. The *net economic impact* of a policy is the present value of consumption under the policy minus the present value of consumption in the base case, in which carbon emissions are uncontrolled. Because the DICE model and its variants incorporate environmental damages into the model, it is possible to calculate an optimal carbon policy, in which the marginal costs of carbon abatement equal the marginal benefits of lower emissions. Typically, simulations involving these models find that the optimal policy is to go slow. Since carbon emissions remain in the atmosphere for several hundred years, the marginal damages resulting from any new emissions are modest. Thus, gradually phasing in carbon reduction lowers the opportunity cost of reducing emissions without having much impact on the global climate.

In the base case of the ENTICE model, ITC increases net economic welfare by 8.3% over a model with exogenous technological change. However, environmental benefits are negligible. Even after 100 years, there is practically no change (less than 1%) in average global temperature when ITC is added

¹⁶ The fossil fuel sector is similar to that used in the RICE model, which is the regional version of DICE. Details can be found in Popp (2003b).

to the model. Nor is there much change in the optimal level of the carbon tax. Rather, welfare benefits come from cost savings, in that emissions goals can be reached more cheaply.

Also notable is that the cumulative effect of R&D is important. Figure 3 compares net welfare with endogenous and exogenous technological change. The figure shows the cumulative welfare gains over time for simulations without a backstop technology. Note that, in the years immediately after the implementation of the carbon tax, the welfare gains are actually lower when ITC is included in the model. Although R&D responds quickly to the carbon tax, in these early years, the opportunity costs of performing additional R&D outweigh the potential benefits of more efficient energy technology. However, since the benefits of R&D are cumulative, future research builds on these early successes, and by 2125, the welfare gains from the model with ITC surpass those in the exogenous technology simulations. As noted in Lesson #2 above, the productivity of future research depends on the productivity of the knowledge on which they can build. Thus, research done today not only offers the benefit of potential energy savings, but also serves to make future research more productive.

Finally, note that the difference between the exogenous and ITC case soon stabilizes. Here we see the importance of diminishing returns over time. Over time, finding new research successes in the energy savings field becomes more and more difficult, so that the welfare gains from ITC cannot grow continuously.

Implication #2 – The opportunity cost of R&D is important

Recall from Lesson #4 that the social returns of R&D are typically higher than the private returns. This offers an important limitation to the effect of ITC. To the extent that energy R&D crowds out other forms of research, similarly productive research in other areas is reduced.

Based on the results from lesson #4 above, in the ENTICE model, I assume energy R&D spending has a social rate of return four times greater than its private rate of return. Fifty percent crowding out is assumed – that is, half of new energy R&D spending displaces other R&D spending, which also has a social rate of return four times greater than its private rate of return. The rate of

crowding out is based on changes in U.S. private sector R&D that occurred between 1972-1998. Sensitivity analysis shows that the opportunity cost is an important limitation on the potential benefits of ITC. When crowding out is removed from the ENTICE model, welfare gains from ITC increase to as much as 43.6%, compared to 8.3% in the base case with partial crowding out. Similarly, with complete crowding out, the welfare gain from ITC falls to as little as 2.1%.¹⁷ Such results help to explain the differences between models. For example, Nordhaus (2002), who finds little impact for ITC, assumes a fixed supply of R&D labor, so that complete crowding out of other R&D occurs as energy R&D increases. In comparison, Buonanno *et al.* (2003) assume that environmental R&D and other R&D are complements, rather than substitutes. Not surprisingly, they find a much larger potential impact for ITC. In addition, these results explain why top-down models assuming a learning curve, but not including policy-induced innovation, tend to find larger benefits from technological change than bottom-down models. Although learning by doing models do include an opportunity cost from increasing investment in alternative energy sources, this cost is simply the loss of one dollar of potential investment elsewhere, as opposed to the loss of more productive R&D dollars from elsewhere in the economy. As a result, learning-by-doing models will overstate the potential gains from technological improvements.

An important policy implication of crowding out is that it limits the potential of ITC under more restrictive climate policies. Since more restrictive policies induce more energy R&D than the optimal policy derived in ENTICE, the opportunity cost of lost R&D in other sectors becomes greater. To see this, consider the results of a simulation restricting global carbon emissions at 1995 levels, taken from Popp (2003b). Whereas including ITC improved welfare by 8.3% in the optimal policy simulation, welfare only improves by 5.6% in this simulation when ITC is included. The importance of crowding out for this result can be seen by eliminating crowding out from the simulation. Doing so increases the net welfare

¹⁷ In Popp (2003b), I present upper and lower bounds of the influence of opportunity cost in the ENTICE model. The upper bound estimates assume that R&D levels change as the opportunity costs change. The lower bound estimates change the opportunity cost, but constrain R&D to remain at the base case levels. The latter can be interpreted as the gains from removing the *assumption* of partial crowding out. They are not the gains that would result if the government intervened to remedy the problem of partial crowding out.

gain by 46%. A similar experiment in the optimal policy scenario only increases the welfare gain by just 14%.¹⁸

Implication #3 – Market failures for knowledge are important

As noted in lesson #4, the high rates of return that make the opportunity cost of R&D important are the result of market failures in markets for knowledge. The public goods nature of knowledge causes firms to underinvest in R&D. To model these imperfections, the base ENTICE model constrains the social rate of return on R&D to be four times greater than that of the return on other investment. The problem of underinvestment could be addressed by subsidies to energy R&D, if government investments in R&D are set so that all social returns are captured. Models that do not include market failures, such as Buonanno *et al.* (2003) implicitly assume that government R&D subsidies sufficient to correct all market failures are in place.

Simulations removing the constraint on returns to R&D suggest that the returns on such subsidies could be quite significant. The gain from ETC for the optimal policy improves to 14.0% when energy subsidies are considered. Similarly, adding R&D subsidies to the endogenous R&D case increases the net welfare gain to seven percent in the more restrictive policy case.

V. Future Research – Whither Diffusion?

The simulation results presented above, while preliminary, not only suggest important policy implications, but also fruitful avenues for future research. In particular, one area of great importance is the diffusion of new energy technologies, particularly across countries. In nearly all of the models discussed above, diffusion is not an issue.¹⁹ For example, both R&DICE and ENTICE are global models. As such, once a technology is available, it can be used worldwide. Similarly, Goulder and Schneider's model focuses only on the United States. Assuming quick diffusion of technologies within a country is

¹⁸ These results use the lower bound of the opportunity costs simulations. This holds R&D the same in both the high and low opportunity cost case, to isolate the opportunity cost of the induced research in the base case.

¹⁹ Buonanno *et al.* (2003) is an exception.

reasonable, and supported by empirical evidence (see, for example, Popp 2001). However, there is no reason to expect that diffusion between countries should occur at the same rate. For example, research at the domestic level might be needed to adopt foreign knowledge to domestic conditions. Just as the empirical literature on knowledge spillovers finds that spillovers are greatest among similar industries, we would also expect adoption of technologies to be quickest in similar countries.²⁰

At the international level, there is a large body of research on technological diffusion in general (surveyed by Keller 2001), but little work focusing specifically on environmental technologies. As noted earlier, the one exception is Lanjouw and Mody (1996). Suggesting the importance of diffusion, they find that the majority of environmental technology patents in developing countries come from foreign countries. Moreover, they find that policies in one nation may affect innovation of technologies in a second nation. Looking at motor vehicles, they note that the majority of vehicle air emissions patents granted in the U.S. are from foreign nations, even though the US was the first country to adopt strict emissions standards.

Given the importance of diffusion for bringing new environmental technologies to developing countries, understanding the general lessons from empirical work on international diffusion is important. There are two potential avenues through which foreign knowledge can have an influence in the domestic economy. First, foreign inventions may be adopted directly by domestic firms. Second, rather than being directly adopted by domestic users, foreign knowledge may affect the productivity of domestic R&D. The “blueprints” represented by foreign patents may serve to inspire additional innovation by domestic inventors. Such productivity increases are knowledge spillovers, as the knowledge represented by the foreign patent creates a positive externality – knowledge from abroad that is borrowed by the domestic inventor (Griliches 1979).

²⁰ At the industry level, Cohen and Levinthal (1989) show that firm R&D increases the ability of firms to absorb knowledge spillovers. If such R&D is needed for spillovers to occur, the opportunity cost of this R&D must be considered. Although such considerations are not included in a global model such as ENTICE, they are important for sector-specific computable general equilibrium (CGE) models.

Since most R&D takes place in just a few countries, both types of diffusion are important. For direct adoption, many of the same lessons from diffusion studies within a country apply here as well. In particular, local conditions are an important determinant of the rate of diffusion across countries. Regarding knowledge spillovers, researchers have followed several strategies. At a macro level, several papers, most notably Coe and Helpman (1995), have estimated the effects of domestic and foreign R&D stocks on total factor productivity (TFP) growth. Typically, the foreign R&D stocks are weighted by international trade flows, so that these papers attempt to measure the effect of international trade on knowledge flows. Such papers typically find a positive effect for foreign R&D, although this effect may vary by country. In general, the relative contribution of foreign knowledge is inversely related to economic size and level of development (Keller 2001). At a more micro level, studies looking at the effect of foreign direct investment (FDI) and foreign knowledge spillovers find mixed results. Whether the effect of FDI on growth is positive or negative depends on individual industry and country characteristics, with inward FDI being more successful for developed countries (Keller 2001).

Such results suggest that diffusion may limit the success of induced innovation as developing countries begin to make greater use of fossil fuels. Since nearly all R&D is performed in the world's richest nations, diffusion of newly developed environmental technologies will be necessary if technology is to play a role in mitigating future carbon emissions from the developing world. Research on both the factors that affect the flows of environmentally-friendly technologies across countries, as well as research on the speed at which such knowledge diffuses, will greatly improve our ability to model the potential effects of technology on the climate change problem. Patents are likely to be a useful tool for such research. Since patents are filed in multiple countries, patent data makes it possible to track the usage of inventions across countries. Such data should be an invaluable source as studies of diffusion progress.

References

- Ahmad, Syed (1966), "On the Theory of Induced Innovation," *Economic Journal*, 76(302), 344-57.
- Bellas, Allen S. (1998), "Empirical evidence of advances in scrubber technology," *Resource and Energy Economics*, 20(4), 327-343.
- Berndt, Ernst, Charles Kolstad, and Jong-Kun Lee (1993), "Measuring the Energy Efficiency and Productivity Impacts of Embodied Technical Change," *The Energy Journal*, 14, 33-55.
- Binswanger, Hans P. (1978a), "The Microeconomics of Induced Technical Change," in by Hans P. Binswanger and Vernon W. Ruttan, eds., *Induced Innovation: Technology, Institutions, and Development*, Baltimore, MD: John Hopkins University Press, 91-127
- Binswanger, Hans P. (1978b), "Issues in Modeling Induced Technical Change," in by Hans P. Binswanger and Vernon W. Ruttan, eds., *Induced Innovation: Technology, Institutions, and Development*, Baltimore, MD: John Hopkins University Press, 128-63.
- Binswanger, Hans P. (1974), "A Microeconomic Approach to Innovation," *Economic Journal*, 84(336), 940-58.
- Caballero, Ricardo J. and Adam B. Jaffe (1993), "How High are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth," in Olivier J. Blanchard and Stanley Fischer, eds., *NBER Macroeconomics Annual 1993*, Cambridge, MA: MIT Press, 15-74.
- Carlson, Curtis, Dallas Burtraw, Maureen Cropper, and Karen Palmer (2000), "Sulfur Dioxide Control by Electric Utilities: What Are the Gains from Trade?" *Journal of Political Economy*, 108(6), 1292-1326.
- Coe, David T. and Elhanan Helpman, (1995), "International R&D Spillovers," *European Economic Review*, 39, 859-887.
- Cohen, Wesley M. and Daniel A. Levinthal (1989), "Innovation and Learning: The Two Faces of R&D," *The Economic Journal*, 99, 569-596.
- David, Paul A. (1969), *Contribution to the Theory of Diffusion*, Stanford Center for Research in Economic Growth, Memorandum #71, Stanford University.
- Davies, Stephen W. (1979), *The Diffusion of Process Innovations*, Cambridge University Press, Cambridge, U.K.
- Ellerman, A. Denny, Richard Schmalensee, Paul L. Joskow, Juan Pablo Montero, and Elizabeth M. Bailey (1997), *Emissions Trading Under the U.S. Acid Rain Program: Evaluation of the Compliance Costs and Allowance Market Performance*, MIT Center for Energy and Environmental Policy Research, Cambridge, MA.
- Evenson, Robert E. and Yoav Kislev (1975), *Agricultural Research and Productivity*, Yale University Press, New Haven, CT.
- Fisher, Carolyn, Ian W. H. Parry, William A. Pizer (1998), "Instrument Choice for Environmental Protection When Environmental Protection is Endogenous," *Resources for the Future Discussion Paper 99-04*.
- Goulder, Lawrence H. and Stephen H. Schneider (1999), "Induced Technological Change and the Attractiveness of CO₂ Emissions Abatement," *Resource and Energy Economics*, 21, 211-253.
- Gray, Wayne B. and Ronald J. Shadbegian (1998), "Environmental Regulation, Investment Timing, and Technology Choice," *The Journal of Industrial Economics*, 46(2), 235-256.

- Griliches, Zvi (1990), "Patent Statistics As Economic Indicators: A Survey," *Journal of Economic Literature*, 28(4), 1661-707.
- Griliches, Zvi (1979), "Issues in Assessing the Contribution of Research and Development to Productivity Growth," *Bell Journal of Economics*, 10, 92-116.
- Griliches, Zvi (1957), "Hybrid Corn: An Exploration of the Economics of Technological Change," *Econometrica*, 25, 501-522.
- Hall, Bronwyn (1995), "The Private and Social Returns to Research and Development," in *Technology, R&D, and the Economy*, Bruce Smith and Claude Barfield, eds., Brookings, 140-183.
- Hicks, John R. (1932), *The Theory of Wages*, London: Macmillan.
- Jaffe, Adam B. (1986), "Technological Opportunity and Spillover of R&D: Evidence from Firms' Patents, Profits, and Market Value," *American Economic Review*, 76, 984-1001.
- Jaffe, Adam B., Michael S. Fogarty, and Bruce A. Banks (1998), "Evidence from Patents and Patent Citations on the Impact of NASA and Other Federal Labs on Commercial Innovation," *Journal of Industrial Economics*, 46, 183-205.
- Jaffe, Adam B., Rebecca Henderson, and Manuel Trajtenberg (1993), "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," *Quarterly Journal of Economics*, 108(3), 577-98.
- Jaffe, B., Richard G. Newell, and Robert N. Stavins (2003), "Technological change and the environment," in K.G.Mäler & J. Vincent (Eds.), *Handbook of Environmental Economics*, (pp. 461-516). Amsterdam: North Holland.
- Jaffe, Adam B., Richard G. Newell, and Robert N. Stavins (1999), "Energy-Efficient Technologies and Climate Change Policies: Issues and Evidence," Climate Issue Brief No. 19. Washington, D.C.: Resources for the Future. Reprinted in Toman, Michael A., ed. (2001), *Climate Change Economics and Policy*, pp. 171-181, Washington, D.C.: Resources for the Future.
- Jaffe, Adam B. and Karen Palmer (1997), "Environmental Regulation and Innovation: A Panel Data Study," *Review of Economics and Statistics*, 79(4), 610-19.
- Jaffe, Adam B. and Robert N. Stavins (1995), "Dynamic Incentives of Environmental Regulation: The Effects of Alternative Policy Instruments on Technology Diffusion," *Journal of Environmental Economics and Management*, 29, S43-S63.
- Jaffe, Adam B. and Maneul Trajtenberg (1996), "Flows of Knowledge From Universities and Federal Labs: Modeling the Flow of Patent Citations Over Time and Across Institutional and Geographic Boundaries," *Proceedings of the National Academy of Sciences*, 93, 12671-77.
- Johnson, Daniel K.N. and David Popp (2003), "Forced Out of the Closet: The Impact of the American Inventors Protection Act on the Timing of Patent Disclosure," *RAND Journal of Economics*, 34(1), 96-112.
- Jorgenson, Dale W. and Barbara M. Fraumeni (1981), "Relative Prices on Technical Change," in *Modeling and Measuring Natural Resource Substitution*, Ernst R. Berndt & Barry C. Field eds., MIT Press, Cambridge, MA, 17-47.
- Keller, Wolfgang (2001), "International Technology Diffusion," *NBER Working Paper #8573*.
- Kerr, Suzi and Richard Newell (2001), "Policy-Induced Technology Adoption: Evidence from the U.S. Lead Phasedown," *Resources for the Future Discussion Paper 01-14*, forthcoming in *Journal of Industrial Economics*.

- Lanjouw, Jean O. and Ashoka Mody (1996), "Innovation and the International Diffusion of Environmentally Responsive Technology," *Research Policy*, 25, 549-571.
- Lanjouw, Jean O., Ariel Pakes and Jonathan Putnam (1998), "How to Count Patents and Value Intellectual Property: Uses of Patent Renewal and Application Data," *The Journal of Industrial Economics*, 46(4), 405-433.
- Lanjouw, Jean O. and Mark Shankerman (1999), "The Quality of Ideas: Measuring Innovation with Multiple Indicators," *NBER Working Paper #7345*.
- Levin, Richard C., Alvin K. Klevorick, Richard R. Nelson, and Sidney G. Winter (1987), "Appropriating the Returns from Industrial Research and Development," *Brookings Papers on Economic Activity*, 3, 783-820.
- Kamien, Morton I. and Schwartz, Nancy L. (1968), "Optimal Induced Technical Change," *Econometrica*, 36(1), 1-17.
- Magat, Wesley A. (1978), "Pollution Control and Technological Advance: A Dynamic Model of the Firm," *Journal of Environmental Economics and Management*, 5, 1-25.
- Mansfield, Edwin (1996), "Microeconomic Policy and Technological Change," in *Technology and Growth: Conference Proceedings*, Jeffrey C. Fuhrer and Jane Sneddon Little, eds., Federal Reserve Bank of Boston, 183-200.
- Mansfield, Edwin (1977), "Social and Private Rates of Return from Industrial Innovations," *Quarterly Journal of Economics*, 91, 221-240.
- Mansfield, Edwin (1968), *Industrial Research and Technological Innovation*, W.W. Norton, New York.
- Milliman, Scott R. and Raymond Prince (1989), "Firm Incentives to Promote Technological Change in Pollution Control," *Journal of Environmental Economics and Management*, 17, 247-165.
- Mountain, Dean C., Bill P. Stipdonk, and Cathy J. Warren (1989), "Technological Innovation and a Changing Energy Mix -- A Parametric and Flexible Approach to Modeling Ontario Manufacturing," *The Energy Journal*, 10, 139-158.
- National Science Board (2002), *Science and Engineering Indicators – 2002*, National Science Foundation, Arlington, VA (NSB-02-1).
- Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins (1999), "The Induced Innovation Hypothesis and Energy-Saving Technological Change," *Quarterly Journal of Economics*, 114, 941-975.
- Nordhaus, William (2002), "Modeling Induced Innovation in Climate-Change Policy," in *Technological Change and the Environment*, Arnulf Grübler, Nebojsa Nakicenovic, and William D. Nordhaus eds., Resources for the Future, Washington, DC.
- Nordhaus, William (1994), *Managing the Global Commons: The Economics of the Greenhouse Effect*, MIT Press, Cambridge, MA.
- Oster, Sharon (1982), "The Diffusion of Innovation Among Steel Firms: The Basic Oxygen Furnace," *The Bell Journal of Economics*, 13(1), 45-56.
- Pakes, Ariel (1985), "On Patents, R&D, and the Stock Market Rate of Return," *Journal of Political Economy*, 93, 390-409.
- Popp, David (2003a), "Pollution Control Innovations and the Clean Air Act of 1990," *Journal of Policy Analysis and Management*, 22(4), 641-660.
- Popp, David (2003b), "ENTICE: Endogenous Technological Change in the DICE Model of Global Warming," *NBER Working Paper #9762*.

Popp, David (2002), "Induced Innovation and Energy Prices," *American Economic Review*, 92(1), 160-180.

Popp, David (2001) "The Effect of New Technology on Energy Consumption," *Resource and Energy Economics*, 23, 215-239.

Rose, Nancy L. and Paul L. Joskow (1990), "The Diffusion of New Technologies: Evidence From the Electric Utility Industry," *RAND Journal of Economics*, 21(3), 354-373.

Schumpeter, Joseph (1942), *Capitalism, Socialism and Democracy*, Harper, New York.

Sternier, Thomas (1990), "Energy Efficiency and Capital Embodied Technical Change: The Case of Mexican Cement Manufacturing," *The Energy Journal*, 11, 155-167.

Figure 1 – The S-shaped diffusion curve







Figure 2 – Induced Innovation and Energy Prices



Source: Popp (2002)

Figure 3 – Welfare gains over time in the ENTICE model



Source: Popp (2003b)

The figure shows the cumulative gains in welfare from an optimal carbon tax. Note that welfare in the induced innovation is initially lower than exogenous R&D, but that induced innovation leads to larger long run improvements than in a simulation with exogenous technological change.

Table 1 – Induced Innovation and the Elasticity of Energy Consumption

	(1)	(2)	(3)	(4)	(5)
Industry	elasticity of energy use with respect to induced innovation	elasticity of energy use with respect to price	total elasticity of energy use with respect to price	% of elasticity due to induced innovation	price elasticity without knowledge stocks
Chemicals	-0.319	-1.378	-1.697	19%	-1.652
Copper	-0.018	-0.235	-0.253	7%	-0.220
Electrometallurgical	-1.827	-1.596	-3.423	53%	-1.478
Pulp & Paper	-0.037	-0.250	-0.287	13%	-0.307

Source: Popp (2001)

The table shows the short-run breakdown in the effects on energy consumption from a change in energy prices. The elasticity of energy use with respect to induced innovation is the percent change in energy consumption resulting from the new technologies induced by a one percent change in energy prices. The elasticity of energy use with respect to price is the change in energy consumption resulting from factor substitution. Column three shows the total effect of a change in energy prices. It is the sum of the elasticities in columns one and two. Column 4 presents the percentage of the total elasticity that is due to induced innovation. Finally, column 5 gives the elasticity of energy with respect to price from a regression that excludes the knowledge stocks, so that only a time trend is used to capture technological change. Note that these elasticities are comparable to the combined elasticities presented in column 3.