

Induced Innovation and Energy Prices

David Popp
The University of Kansas
August 31, 1999

Abstract

This paper uses U.S. patent data from 1970 to 1994 to estimate the effect of energy prices on energy-efficient innovations. It focuses on both demand-side factors, which spur innovative activity by increasing the value of new innovations, and supply-side factors, such as scientific advancements that make new innovations possible. Using patent citations as a measure of the usefulness of the existing base of scientific knowledge, I construct *productivity estimates*, which measure the usefulness of the existing stock of knowledge to inventors in a given energy field for any given year. These estimates are then combined with data on demand-side factors to estimate a model of induced innovation in energy technologies. The results indicate that both energy prices and the supply of knowledge have strongly significant positive effects on innovation. Furthermore, I show that omitting the productivity estimates biases the estimation results. The paper concludes with a discussion of the implication of this work for environmental policy.

The author would like to thank William Nordhaus, Joel Waldfoegel, Robert Evenson, Ariel Pakes, Dietrich Earnhart, Tom Weiss, and two anonymous referees for helpful comments. Despite their useful suggestions, however, all mistakes are solely the responsibility of the author. In addition, the assistance of Adam Jaffe, who provided the programming code for the patent citation regressions, is gratefully acknowledged. Finally, thanks to Robert Evenson, Jon Putnam, and Samuel Kortum for their help in obtaining patent data. Partial financial support for this paper comes from DOE grant 593A3140217.

The development of new technologies can be of great value to any society. As such, it is desirable to understand the process by which it occurs. In particular, we want to know what determines the pace and timing of technological change and especially whether the process can be influenced by policy. One area for which understanding the link between policy and technological change is particularly important is the environment. Many environmental problems, such as global warming, are long-term problems that technological progress may play a crucial role in ameliorating. However, there has been little empirical evidence about policy-induced development of environmentally-friendly technology, such as the relationship between prices (or other policy measures) and technological change.¹ As a result, most models of environmental problems treat technological change as exogenous, and as a consequence cannot examine potentially important links between policy and technological change.² Furthermore, models that do incorporate endogenous technological change have been hampered by the lack of empirical evidence on the links between environmental policy and innovation.³

This paper addresses the shortcoming in our understanding of endogenous technological change by using U.S. patent data from 1970 to 1994 to study the effect of energy prices on energy-efficient innovations. Although its main focus is the influence of energy prices on innovative activity related to energy-efficiency, it also sheds light on the more general question: *what factors influence inventive activity?* In 1932, J.R. Hicks introduced the theory of induced

¹ The existing literature on induced environmental innovation includes Lanjouw and Mody (1996), Jaffe and Palmer (1997), and Newell, Jaffe, and Stavins (1998). However, only Jaffe and Palmer include elasticities that could be incorporated into large-scale models of the environment.

² Two recent exceptions are Goulder and Schneider (1999) and Nordhaus (1997). As evidence that there is still much work to be done in correctly modeling induced technological change, the two papers reach opposite conclusions about the effect that innovation induced by a carbon tax may have on reducing emissions.

³ For example, Goulder and Schneider (1999) state that their “ability to generate more precise estimates [of the contribution of endogenous technological change] is fundamentally limited by the absence of empirical estimates on the relationship between R&D expenditure and technological change.”

innovation, which states that changes in relative factor prices should lead to innovations that economize the relatively expensive factor. Interest in the microeconomics of induced innovation increased after papers by Ahmad (1966), Kamien and Schwartz (1968), and Binswanger (1974, 1978a, 1978b). However, empirical testing of the induced innovation hypothesis was hampered by a lack of available data. The U.S. patent data used in this paper allows me to overcome this deficiency.

Because the induced innovation literature treats as exogenous the existing stock of knowledge on which inventors can build, it ignores the determinants of that base of knowledge. As such, the induced innovation literature cannot truly endogenize the path of technological change (Nordhaus 1973). In particular, induced innovation theories do not capture the effect of current research on future research efforts. For example, if diminishing returns to research exist, increases in the current level of R&D may make future R&D more difficult. In these instances, the induced innovation literature cannot identify the tradeoff between current research on a factor and the productivity of future research on the factor. In contrast, *technology-push* theories of R&D emphasize the importance of technological opportunity to innovation, capturing the endogenous aspects of the process.⁴ In this paper, I make use of patent citation data to measure technological opportunity, so that I may combine information on both market demand and technological opportunity in an empirical study of the determinants of innovation. I show that including information on both is crucial for accurate estimation of the factors that drive innovation.

The paper begins by detailing the construction of the energy patent data set and discusses how it will be used in the regressions that follow. However, to estimate the elasticity of energy

⁴ For examples see Scherer (1965), Schmookler (1966), Rosenberg (1982, 1984), and Mowrey and Rosenberg (1979).

patents with respect to energy prices, more information is needed. A measure of the usefulness of the existing base of scientific knowledge must also be included. Using patent citations, I construct *productivity estimates* that capture changes over time in the usefulness of energy patents for future inventors. This is done in section II. The productivity estimates are used in section III, along with the rest of the data in section I, to test the importance of both demand and supply side influences in the development of new energy technologies. The results indicate that both energy prices and the expected productivity of R&D have strongly significant positive effects on innovation. Furthermore, I show that omitting the productivity estimates leads to underestimating the true elasticity between prices and patents. I conclude by discussing the implication of these results for both environmental policy and the economics of technological change.

I. Modeling & Data

Despite the theoretical advances of the induced innovation literature during the 1960's and 1970's, sufficient data to test the hypotheses were unavailable. As a result, most existing tests of the induced innovation hypothesis have been indirect tests focusing on the results of innovation rather than the innovation process itself. One of the major contributions of this paper is the creation of a data set of energy patents in the United States from 1970-1994, made possible by the computerization of patent data that has occurred. These improved data sources permit me to develop a direct test of the induced innovation hypothesis.

A. The Energy Patent Data Set

When a patent is granted in the United States, it is given a U.S. classification number. There are over 300 main classification groups and over 50,000 subclassifications. The first step in assembling the data set for this paper was to identify subclassifications pertaining to energy efficiency. Using resources from the Department of Energy and from the academic sciences, I

identified several areas of research in the energy field. Descriptions of these technologies were matched with U.S. patent subclassifications. Technologies for which no clear subclassification exists were eliminated. The resulting set of subclassifications was then sorted into 22 distinct technology groups to be used in this paper. The technology groups include 11 groups pertaining to energy supply, such as solar energy, and 11 relating to energy demand, such as methods of reusing industrial waste heat. Table 1 lists the technology groups included in this paper.⁵

Using data from the MicroPatent CD-ROM database of patent abstracts and additional data from the U.S. Patent and Trademark Office, I identified all patents in the 22 technology groups that were granted in the U.S. between 1970 and 1994.⁶ For the purposes of this paper, only patents granted to Americans are included, since foreign inventors are likely to be influenced by factors not included in the data described below. Also, to identify the effect that federal R&D spending has on private research efforts, patents held by government agencies are not included.

For each technology group, patents are sorted by the year of application.⁷ In the United States, information on patents is not made public until the patent is granted. Thus, only *successful* patent applications are included in the data set. Several papers have found that patents, grouped by the date of application, are a good indicator of R&D activity (see Griliches (1990) for a

⁵ Interested readers may download a more thorough description of the technologies chosen at <http://eagle.cc.ukans.edu/~dpopp/dpopp.html>.

⁶ The MicroPatent database contains every U.S. patent issued from 1975-1994. To consider data from the first energy crisis of 1973, additional patent data was obtained from the Classification and Search Support System (CASSIS), available from the U.S. Patent and Trademark Office, so that the data set in this paper extends backwards to 1970. Unfortunately, for the patents obtained from CASSIS, some of the data found in the MicroPatent database, such as patent citations, are not available.

⁷ Note that the patent application data is taken from the front page of the granted patent, and does not include the date of application of earlier continuations and divisions that may exist. A check of a random sample of patents in the data set indicates that about 20 percent of patents are affected. These patents would have earlier application dates than is given on the front page. Since each year contains patents that were erroneously assigned to it, rather than an earlier year, *and* omits patent applications that were filed but then continued in a later year, most of the error should cancel out. Any remaining error would bias the estimated elasticities downward slightly, since more incorrectly assigned patents would occur in years of heavy patenting activity, such as when energy prices are high.

survey). Since a patent application is only made public when a patent is granted, the data have been scaled up to account for patents applied for but not yet granted. Using the distribution of the lag between application and grant for all energy patents in the data set, the percentage of patents remaining to be granted after 1994 was determined for each application year and added to the data. Because a large number of patents applied for in recent years have yet to be granted, and since recently granted patents have not had a chance to be cited yet, only patent applications through 1990 are used. Given the distribution of lags between application and grant in the data, just 1 percent of patents applied for in 1990 remain to be granted.

Table 2 provides the annual patent count in each of the technology groups from 1970 to 1993⁸. Figure 1 illustrates trends in the data for five of the technology groups, with 1981 normalized to 100.⁹ Included for reference in these figures is an index of the cost of energy in dollars per million British thermal units (Btu). For most of the technology groups, there is a jump in patent applications during the energy crises of the 1970's, suggesting that energy prices do play an important role in inducing energy-efficient technological change.¹⁰ This relationship is further evident in Table 3, which presents the correlation between patent applications and energy prices. Note that the strongest correlations are with current energy prices, and that the effect of lagged energy prices drops quickly.

In addition to illustrating the strong effect that energy prices have on patenting activity, Figure 1 also highlights the importance of considering the returns to R&D. Note that energy

As I nonetheless find significant coefficients in the regressions, the effect of this error appears to be small.

⁸ For informational purposes, the figures include patent applications through 1993. However, because many of these patents remained to be granted by 1994, only patent applications through 1990 are used in the subsequent regressions.

⁹ Additional figures are available from the author by request.

¹⁰ Notable exceptions include fuel cells, the use of waste as fuel, and continuous casting. Reasons why specific technologies may differ are discussed in the next section.

prices do not peak until 1981. Nonetheless, patenting activity in most of these technology groups reaches its peak in the late 1970's. Had the returns to energy R&D remained constant over time, we would expect patenting activity in these fields to remain high until prices began to fall. That patenting activity drops before prices do suggests that the possibility of diminishing returns to research should be explored.¹¹ The use of patent citation data makes that possible.

B. Other Data

In addition to data on energy patents, the following other data are used in the paper. Data on energy prices are taken from the *State Energy Price and Expenditure Report*, published by the Energy Information Administration. Prices are in constant 1987 dollars per million British thermal unit (Btu), deflated by a GNP fixed-weight price deflator. Prices are available for different sectors of the economy, such as industrial and residential, and by type of fuel, such as coal or fossil fuels. In addition, an index of total energy prices is provided. Data for government R&D are taken from Public Citizen (1992), which collected data on federal R&D expenditures by energy technology from the annual U.S. government budget publications

I also include data for several technology-specific variables. These variables capture the effects of characteristics unique to the individual technology groups, and help to explain some of the individual trends in the data described in footnote 10. For continuous casting, the data set includes the price of ore purchased by steel producers. Since continuous casting makes more efficient use of raw materials, there is a positive relationship between the price of ore and continuous casting patents. For the use of waste products as energy, the price of waste to the utilities is included in the data set. This captures the increased supply of waste available as fuel

¹¹ One might argue that another possible explanation for the early decline in energy patents is that political support for energy R&D may have changed. However, the drop in patenting occurs during the Carter years, when

due to concerns over declining landfill space during the 1980's. Finally, for fuel cells, the excess capacity of electric utilities is included. In addition to using fossil fuels more efficiently, fuel cells also offer the advantage of modularity. Individual cells are small, and can be linked to generate as much power as needed. They can be installed in small spaces, and little lead-time is needed to set up a plant that uses fuel cells to generate electricity. Having over-estimated future electricity demand in the 1970's, the electric utility industry built too many new, large power plants. In this climate, fuel cells offered a more flexible alternative when increased power-generating capacity was needed in the 1980's. (OTA, 1991) We would expect a positive correlation between excess capacity and fuel cell patents.

C. Modeling

The data described above will be used to test the induced innovation hypothesis for energy research. Because much of the research and development process is poorly understood, specifying a structural model of the determinants of energy patents is difficult. Rather, I use a simple log-log form regression of successful patent applications on energy prices, the productivity estimates, and the other variables discussed in section *IB*. Such a specification allows the resulting coefficients to be interpreted as elasticities, which can then be used by modelers of environmental problems such as global warming.

Using patent data as a measure of research output poses several complications. The first is that the propensity to patent varies widely by industry. In some industries, such as the chemical industry, many new innovations are patented. In other industries, secrecy is a more important means of protection. In these industries, the cost of revealing an idea to competitors is often not worth the gains from patent protection. As a result, the correlation between R&D and patents

support for energy research was at its highest.

varies across industries.¹² However, our interest lies more in the time series aspects of the data. It does not matter whether there are twice as many patents in one field as another. What does matter is whether the number of patents in each field increases with an increase in energy prices. As long as the tendency to patent remains the same in each field across time, variations in the tendency to patent across different fields will not pose a problem.

Unfortunately, the second difficulty with using patent data is that it is not clear that the tendency to patent is the same across time. Over time, the ratio of patents to R&D expenditures has fallen in the United States (as well as in other industrialized nations). Some researchers, most notably Evenson (1991), consider the falling ratio to be evidence of diminishing returns to R&D.¹³ Similarly, Kortum and Lerner (1998) argue that a recent upswing in patenting activity in the United States is due to the increasing fertility of new research opportunities. Other researchers, most notably Griliches (1989), claim that research opportunities have not declined.

¹² Levin *et al.* (1987) discusses the variation in patenting behavior across industries.

¹³ In this paper, diminishing returns to research refers to the expected return on the *inputs* to the research process, not the returns to the input. The notion that there are increasing returns to the *output* of knowledge, usually attributed to the public good nature of knowledge, is by no means compromised by claiming that the *inputs* to research experience diminishing returns. Diminishing returns to research simply implies that it becomes more and more difficult to develop new inventions as time progresses.

Griliches argues that the fall in the patent-to-R&D ratio is due to changes in the willingness of inventors to patent new inventions. An exogenous fall in the willingness to patent – caused, for example by changes in patent laws that affect the benefits of holding a patent – would result in a falling patent-to-R&D ratio even if the productivity of research spending remained the same.

The challenge raised by the falling patent/R&D ratio is to determine the cause of the decline. It could stem either from changes in the perceived benefits of patents or from changes in the probability of success of R&D. This paper takes two steps to help identify and control for the cause. First, the dependent variable is the *percentage of successful domestic patent applications per year that are in each technology field*. Using the percentage of applications, rather than a raw count of applications, accounts for growth in the economy and exogenous changes in patenting behavior. Changes that affect *all* patent classifications will lead to a change in the total number of patent applications in a given year. Thus, by using a percentage of patent applications, exogenous changes in the propensity to patent are removed from the data.¹⁴

The second step to identifying changes in patenting behavior over time is to include a measure of the marginal productivity of R&D. In section II, I use patent citations to construct such a measure. As section II will show, the estimates of the marginal productivity of R&D suggest that there are diminishing returns to R&D over time. Since I control for exogenous changes in the propensity to patent by using a percentage of patent applications as the dependent variable, a significant positive coefficient on the lagged productivity estimates suggests that diminishing returns to research explain at least part of the decline in the patent/R&D ratio.

¹⁴ During the years studied in this paper, total domestic patent applications start at 54,894 in 1971, fall to a low of 31,548 in 1985, and rise to 59,032 by 1990.

The model to be estimated is as follows. Let $EPAT_{i,t}$ represent the number of successful non-government American patent applications for technology field i in year t . $TOTPAT_t$ represents the total number of successful non-government American patent applications in year t . $P_{E,t}$ is the price of energy in year t , and $P_{E,t-1}$ is the price of energy for the previous year.¹⁵ $\mathbf{j}_{R,i,t}$ represents the marginal productivity of research inputs. \mathbf{Z} is a vector of the other independent variables described in section IB, such as R&D spending by the U.S. Department of Energy. For any energy-saving technology, i , I estimate a model of the form:

$$(1) \quad \log\left(\frac{EPAT_{i,t}}{TOTPAT_t}\right) = \mathbf{m}_i + \mathbf{l}_i \log P_{E,t} + \mathbf{r} \log P_{E,t-1} + \mathbf{q}_i \log \mathbf{j}_{R,i,t} + \boldsymbol{\eta}_i \log \mathbf{Z} + \mathbf{e}_{it},$$

$$i = 1, \dots, 22; \quad t = 1, \dots, 20.$$

Because decisions on government funding of energy R&D are correlated with energy prices, instrumental variables are used for federal energy R&D spending. Lagged federal energy R&D and a dummy variable representing the lagged political party of the President are used as instruments.¹⁶ Finally, estimates of the marginal productivity of research are needed. In section II, I make use of patent citation data to construct these estimates.

II. Patent Citations and the Existing Stock of Knowledge

When a patent is granted, it contains several citations to earlier patents that are related to the current invention. The citations are placed in the patent after consultation between 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

¹⁵ Ideally, we would prefer to have the expected future price of energy in the regression. Since we do not observe expected prices directly, current and lagged values of energy prices are used instead.

¹⁶ Federal funding for energy R&D fell dramatically a year after President Reagan took office in 1981. A lagged value of the party is used to account for delays in the budgeting process. For example, the fiscal year 1981 budget would have been signed by President Carter before he left office.

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 current 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.¹⁷ For the same reason, inventors have incentives 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.

This paper uses patent citations as evidence of the existing state of technology when the invention was completed. The assumption is that citations indicate a flow of knowledge. Thus, citations to an earlier patent suggest that the patent provided technological opportunity to the new patent. Frequent citations to a patent provide evidence that the knowledge embodied in that invention has been particularly useful to other inventors.¹⁸ We would expect the marginal productivity of research to be greater in the years immediately after useful patents were revealed.

A. Estimating the probability of citation

To begin the formal analysis of patent citations, note that a simple count of subsequent citations is not enough, since the raw number of citations to any patent depends on the total number of patents that follow. Instead, it is necessary to look at the probability of citation. In addition, exogenous factors have changed citing behavior over time. Regression analysis controls

¹⁷ “Narrowing the realm” means that the holder of a patent 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 invention.

¹⁸ Jaffe, Fogarty, and Banks (1998) examine the relationships between knowledge flows and patent citations. Their research includes interviews with scientists, R&D directors, and patent attorneys. They find that, at the level of individual patents, not all citations are indicative of knowledge flows. Other concerns, such as strategically including irrelevant patents to satisfy the patent examiner, add noise to the citation process. However, at more aggregate levels, such as the patents for an organization or firm, they find that patent citations do seem indicative of knowledge flows.

for such factors.

For each technology group, potentially cited patents are sorted by the year in which the patent was granted, which is denoted *CTD* for “year cited.” In the United States, a patent application is not made public unless the patent is granted. Thus, the year of grant is the year in which the patented innovation entered the public domain. The patents that do the citing are sorted by year of application, and are denoted *CTG*, for “citing patent.” Sorting these patents by the year of application allows the results to correspond to the counts of patent applications created in section I.

In principle, regression analysis could be used on each individual patent in the data set. However, the dependent variable – the actual number of citations – is zero for most patents, as most patents most patents are never cited. As a result, the data are sorted by cited and citing year into cohorts of patents that could potentially cite each other. Separate cohorts are constructed for each technology group, *i*. Only citations made by other patents in the technology group are considered.¹⁹ A total of 2,688 *i,CTD,CTG* cohorts are created. For example, one cohort might be citations to all solar energy patents granted in 1975 made by solar energy patents applied for in 1980.²⁰ Denoting citations as *c* and the number of patents from each year as *n*, the probability of citation, *p*, for patents within each cohort is:

$$(2) \quad p_{i,CTD,CTG} = \frac{c_{i,CTD,CTG}}{n_{i,CTD}n_{i,CTG}}.$$

To control for factors that affect the likelihood of citation, I make use of a model first used by Caballero and Jaffe (1993) and Jaffe and Trajtenberg (1996) to estimate the probability of

¹⁹ Similar results are obtained when citations to all patents are considered. These are discussed in Popp (1998).

²⁰ Although the data set includes patents applied for since 1970, citations are only available for the patents granted

a patent being cited by subsequent patents. The model uses an exponential distribution to model flows of knowledge. The equation estimated takes the form:

$$(3) \quad p(i, CTG, CTD) = \mathbf{a}(i, CTG, CTD) \exp[-\mathbf{b}_1(CTG-CTD)][1-\exp(-\mathbf{b}_2(CTG-CTD))].$$

\mathbf{b}_1 represents the rate of decay of knowledge as it becomes obsolete, and \mathbf{b}_2 is the rate at which newly produced knowledge, as represented by a newly-patented innovation, diffuses through society. $\mathbf{a}(i, CTG, CTD)$ are parameters capturing attributes of the citing or cited patents that may influence the probability of citation.²¹ These include:

- the usefulness of the knowledge represented in the patent being cited ($\mathbf{a}_{i,CTD}$),
- the frequency by which patents applied for in the citing year cite earlier patents (\mathbf{a}_{CTG})²²,
and
- the frequency of citations within each technology group (\mathbf{g}).

The \mathbf{a} parameters are based on Jaffe's earlier work. \mathbf{g} is added for this paper. It captures the effect of the size of the technology group. About half of all patent citations are to patents in the same classification (Jaffe, Henderson, and Trajtenberg, 1993). However, the technology groups used in this paper range from groups with one or two subclassifications to groups with patents from many different broad classifications. Technology groups with broad definitions are more likely to include subclasses that are not strongly related. As a result, citations to other

after 1975. Because of the lag between the application and granting of a patent, the average of which is two years, I am able to consider citations made by all patents applied for in 1974 or later.

²¹ Although the model requires us to estimate attributes associated with the cited year, the citing year, and the lag between them, it is possible to identify attributes related to all three because the age of patents enters the model non-linearly.

²² Changes in citing behavior over time must be accounted for because of institutional changes at the patent office that make patents more likely to cite earlier patents than was previously true, even if all other factors are equal. In particular, two changes have played an important role. First, computerization of patent office records has made it easier for both patent examiners and inventors to locate other patents similar to the current invention. Second, increasing legal pressure has made it more important for examiners to be sure that all relevant patents are cited. Since institutional changes will affect all patents equally, this parameter is not indexed by i .

patents in the group are less likely in these groups.

Note that the first parameter, $\mathbf{a}_{i,CTD}$ is the value of interest for this paper. It is $\mathbf{a}_{i,CTD}$ that tells us the likelihood that patents from year CTD will be cited by subsequent patents. The remaining parameters control for other facets of the patenting process that might affect the likelihood of citation. Higher values of $\mathbf{a}_{i,CTD}$ indicate that the patents in question are more likely to be cited by subsequent patents. This implies that the knowledge embodied in those patents is particularly useful. The expected returns on R&D to researchers building on the knowledge in these patents should be larger than for patents from other years. Thus, lagged values of $\mathbf{a}_{i,CTD}$ can be interpreted as a measure of the expected return of R&D. I refer to $\mathbf{a}_{i,CTD}$ as the *productivity parameter*.

Figure 2 provides examples of the trends in citation data for solar energy and fuel cells. The graphs plot the probability of patents granted in year CTD being cited by other patents in the same technology group x -years after the year of grant. Each line represents citations to patents granted in a given year. The x-axis measures the lag in years since the patent was granted, and the y-axis shows the probability of a patent from year CTD being cited by a patent x -years later. The productivity parameter, $\mathbf{a}_{i,CTD}$, can be visualized as the y-intercept in the figures. Patents from productive years, such as solar energy patents from 1975, have higher y-intercepts. Note that the pattern of decay is similar for patents of different vintages, so that productive patents, such as solar energy patents from 1975, remain more likely to be cited than other patents, even after a lag of several years. Also, note that the probability of citation falls over time, suggesting that the

decay of knowledge is more influential than the diffusion of knowledge in determining the probability.²³

Estimation of the values of the productivity parameter, $\mathbf{a}_{i,CTD}$, proceeds as follows. Inserting the variables defined above and adding an error term, $\mathbf{e}_{i,CTD,CTG}$, to equation (3) yields:

$$(4) \quad p_{i,CTD,CTG} = \mathbf{g}_i \mathbf{a}_{i,CTD} \mathbf{a}_{CTG} \exp[-\mathbf{b}_1(CTG-CTD)] \{1 - \exp[-\mathbf{b}_2(CTG-CTD)]\} + \mathbf{e}_{i,CTD,CTG}.$$

Equation (4) is estimated using non-linear least squares, using all patents granted from 1970 to 1989 as the cited years, and all patents applied for from 1974 to 1991 as the citing years.²⁴ Because some of the technology groups have few patents, and thus few citations, the productivity parameter can only be estimated for 12 of the 22 technology fields. To identify the parameters, it is necessary to normalize one of each of the \mathbf{a} 's to be 1.²⁵ Finally, since this is grouped data, the observations are weighted by $\sqrt{(n_{CTD})(n_{CTG})}$ to avoid problems with heteroskedasticity (Greene, 1993).

Figure 3 displays the results of estimation. For each technology included, a plot of the productivity parameter is presented. Numerical results for the remaining parameters are given after the plots. Readers interested in a more detailed discussion of these other parameters are referred to Popp (1998). To interpret the plots of the productivity parameter, note that the productivity parameters in 1970 are normalized to 1. Estimates greater than one for a given year

²³ In these figures, the probability of citation is highest in the year after the patent is granted. This is consistent with the data in Jaffe and Trajtenberg (1996). In their paper, they find that most cites occur in patents *granted* three years after the initial patent. In this data set, citing patents are sorted by *year of application*. Since there is, on average, a two year lag between the initial patent application and the granting of the patent, the results are consistent.

²⁴ The model does not converge when estimating separate \mathbf{a} 's for every possible year. Since the main parameters of interest are the productivity parameters, $\mathbf{a}_{i,CTD}$, separate coefficients are obtained for each of these. The parameters for institutional changes are grouped by two-year periods. Thus, \mathbf{a}_1 represents citation practices in 1974-1975, \mathbf{a}_2 citation practices in 1976-1977, etc.

²⁵ For cited years, \mathbf{a}_{1970} is normalized to 1. For citing years, $\mathbf{a}_{1974-75}$ is normalized to 1. Finally, \mathbf{g} is normalized to

indicate that patents granted in that year are more likely to be cited by future patents than patents which were granted in 1970, and estimates less than one indicate that those patents are less likely to be cited. A downward trend in the productivity parameter is found for most of the technology groups. Even after controlling for exogenous factors such as the number of opportunities for citation, newer patents are less likely to be cited than older patents. This suggests that there are diminishing returns to energy research over time.²⁶

III. Results

Having obtained estimates of the productivity of knowledge in section II, I now move to estimating the induced innovation relationship. Lagged values of the productivity estimates just obtained are used as measures of the marginal productivity of research at time t , making it possible to control for supply-side factors that affect the level of innovation, as well as demand-side factors such as energy prices. Substituting $\mathbf{a}_{i,t-1}$ for $\mathbf{j}_{R,i,t}$ in equation (1), the equation to estimate is:

$$(5) \quad \log\left(\frac{EPAT_{i,t}}{TOTPAT_t}\right) = \mathbf{m}_i + \mathbf{l}_i \log P_{E,t} + \mathbf{r} \log P_{E,t-1} + \mathbf{q}_i \log \mathbf{a}_{i,t-1} + \boldsymbol{\eta}_i \log \mathbf{Z} + \mathbf{e}_{it}$$

$$i = 1, \dots, 17; t = 1, \dots, 20$$

Technology groups that contain few patents are grouped together, resulting in 17 different

I for continuous casting patents.

²⁶ One caveat to interpreting the results of diminishing returns is necessary. This paper focuses on returns to a narrowly defined group of technologies. However, one should not jump to the broader conclusion that returns to R&D must be falling across all technologies. As the expected success rate of energy R&D falls, we would expect R&D efforts to move away from energy and into other, more productive fields. Only if the new fields to which R&D shifts are less productive than previous fields of research would there be diminishing returns to R&D in the economy as a whole. Studies such as this one, applied to other technological fields, could help to shed some light on this question.

technology groups included in the estimation.²⁷ A time trend is used to control for the returns to R&D in technology groups that were too small to obtain a productivity parameter in section II. If the returns to research are diminishing over time, the coefficient on the time trend will be negative.

For concise presentation of results, I pool the technology groups together to obtain single estimates for each parameter.²⁸ In addition, two separate regressions are run pooling only the supply technologies and only the demand technologies. Because the supply technologies are less established than the demand technologies, there are differences in the results across the groupings. Since some of the error in each equation is correlated across technology groups, feasible generalized least squares (FGLS) estimation is used.²⁹ Results are corrected for autocorrelation using the Prais-Winsten transformation, and are presented in Table 4. T-stats for each coefficient are included in parentheses. Almost all the parameters are highly significant. Detailed interpretation of the results follows.

A. The Effect of Energy Prices on Innovation

The results of greatest interest are the effect of energy prices on innovative activity. As Table 4 shows, energy prices play an important role in determining the level of energy innovation. The elasticity of energy patents with respect to current energy prices is 0.125, and

²⁷ Two groups of miscellaneous patents were created: *Other Supply Technologies*, which contains the use of waste gases as fuel, Ocean Thermal Energy Conversion (OTEC), and other uses of natural heat, and *Other Demand Technologies*, which contains carbothermic processing of aluminum, the use of black liquor in paper manufacturing, insulated windows, and compact fluorescent lightbulbs.

²⁸ Estimates for individual technologies, found using seemingly unrelated regression analysis, are available from the author on request.

²⁹ For example, using current and future energy prices is an imperfect substitute for expectations of future prices. Error due to differing expectations (e.g. higher expected prices in 1979 than 1984, although current prices are nearly the same in each year) would affect all the technology groups. In contrast, randomness in the R&D process will affect each technology differently.

with respect to lagged energy prices is 0.462. That patenting activity responds so quickly to prices may seem surprising. However, recall that successful patent *applications* are used as the dependent variable. If the patent application is granted, patent protection is given from the date on which the application was filed (the priority date). Since the costs of applying are low, it is in the inventors interest to file patent applications early in the inventive process, so that an early priority date is established.

Dividing patents into two separate groups according to supply and demand technologies yields more interesting results. The effect of current energy prices is stronger for supply technologies (1.080) than for demand technologies (0.801). However, what is most interesting is that the elasticity with respect to lagged energy prices is negative for the supply technologies. Why might the results differ between the supply and demand technologies? Most likely, these differences are due to the maturity of the technologies. Most of the demand technologies are well-established technologies that were in use both before and after the energy crisis. In contrast, many of the supply technologies were being introduced for the first time. The know-how for many of these technologies may have existed before the energy crisis, but bringing them to market would not have been feasible until energy prices were high. Thus, some of the new patented innovations may simply have been taken “off the shelf” and brought to market when the conditions were right. Little new research would have been needed. In contrast, making improvements to existing technologies, such as most of the demand technologies, would require starting new R&D projects. These results suggest that part of the first wave of innovation after the energy crisis was not due to new ideas being discovered, but rather the introduction of these existing technologically feasible ideas to the marketplace.

B. The Returns to R&D

The regression results show that not only do prices play an important role in inducing new energy innovations, but that the marginal productivity of R&D is also an important factor. The coefficients of the lagged productivity estimates are positive and significant in all three regressions. Furthermore, because there is much variation in the productivity parameters, the magnitude of the effect of productivity is large. The average change in productivity estimates from year to year is 26 percent.³⁰ As a result, changes in the productivity of research would change patenting activity in an average year by 6.8 percent. For comparison, the average change in patenting activity due to a change in prices is about 3.4 percent.³¹ Even during the peak of the energy crisis, energy prices only cause a 5.2 percent increase in patents.³²

Furthermore, it can be shown that omitting the marginal productivity of research from the regressions biases the results. As Figure 1 shows, patenting activity in the energy fields increase quickly when energy prices rise, but begin to fall before energy prices fall. Since the productivity estimates tend to fall over time, the positive coefficients on the productivity estimates, as well as the negative coefficient on the time trend, suggest that diminishing returns to research contribute to the quick fall in energy patenting activity. Omitting these controls for the productivity of research should lead to lower estimates of the effect of prices on patenting activity.

This is confirmed in Table 5A, which omits both the productivity estimates and the time trend from the pooled regressions. As expected, the effect of energy prices falls, especially for the

³⁰ This figure uses the absolute value of changes from year to year.

³¹ The calculation is as follows. The average change in energy prices is only 5.78 percent. Multiplied by the sum of the price elasticities (0.587), we find that the average change in patenting activity due to a change in prices is about 3.4 percent.

³² This figure uses the percent change in energy prices in the years 1974 to 1980, which are the years used in the patent citation analysis.

demand technologies. In the regression including all technologies, current energy prices appear more important, but this change is more than offset by a negative coefficient on lagged energy prices. Also, omitting controls for the productivity of research changes the sign on government R&D spending from positive to negative. For each regression, an F-test rejects the null hypothesis that the coefficients of the productivity estimates and time equal zero.

Next, I ask whether using patent citations to obtain estimates of the marginal productivity of knowledge is necessary, or whether the returns to research can be adequately captured by a time trend. Table 5B provides the results of pooled regressions in which the productivity estimates from section II are replaced by a time trend, so that a time trend is used for all of the technology fields. The results show the value of including the patent citation data. The adjusted R-squares of all three regressions are much lower than the original regressions. Once again, the coefficient on lagged energy prices becomes negative in the overall regression. More astonishingly, in the demand technologies, both price coefficients become negative! The effects of government R&D are also diminished by using a time trend exclusively.

Including the productivity estimates is important because the productivity of knowledge does not monotonically fall over time. In some technology groups, such as fuel cells and continuous casting, the productivity estimates peaked around 1980. In general, these are groups in which there was heavy patenting activity before the sample period, during the 1960's. Diminishing returns led to a low marginal productivity of R&D during the 1970's, until new discoveries came along to once again make research in these fields promising.

The significance of the stock of knowledge on innovation, coupled with evidence of diminishing returns to research within the energy field, have important implications for dynamic models of environmental policy. The price elasticities found above suggest that the reaction of the

research community to a change in policy, such as a carbon tax, will be swift. Higher prices will quickly lead to a shift towards environmentally-friendly innovation. However, since there are diminishing returns to research in a given field, firms will shift their research towards more productive areas of study as the marginal productivity of such research declines. The burst of patenting activity resulting from a policy change is likely to be short-lived.

The results also present evidence that diminishing returns are a factor in the falling patent/R&D ratio. Recall that one of the benefits of using a percentage of successful patent applications as the dependent variable was that it controls for the effect of exogenous factors to the propensity to patent, such as changing patent laws. Even with such exogenous factors controlled for, there is still variation in the data that is explained by the downward trends in the productivity of knowledge.³³ Also, note that many of the productivity estimates show an upswing near the end of the 1980's. This supports the recent results by Kortum and Lerner (1998), who argue that increasingly fertile technology caused a recent surge in patenting in the United States.

C. Federal Energy R&D Spending

Interpreting the effect of government R&D spending on patents is difficult. The focus of federally funded energy R&D changed after President Reagan took office in 1981. Before Reagan took office, federal energy R&D policy included the goal of accelerating the development of new marketable technologies. Support was given to large research projects such as the synfuels program aimed at creating synthetic fuels from coal. In this case, federally funded energy R&D could be a substitute to private innovation. After Reagan's election, government funding for energy research was cut significantly. DOE support for research was limited to long-

³³ This result is examined more thoroughly in Popp (1998).

term, high-risk projects. (Cohen and Noll, 1991). The DOE focused its efforts on the early stages of research and development – basic research to promote general knowledge, and the early stages of applied R&D, designed to test the feasibility of new ideas. It is expected that private firms will continue the R&D process by developing commercially acceptable products. (U.S. Department of Energy, 1987) As such, federally funded R&D should be a compliment to private innovation.³⁴

Although the changes in federal research policy make identifying the effects of research difficult, some broad insights can be taken from the regression results. In general, the results appear to indicate that federal research spending is not a productive means of increasing energy R&D. Only for the supply technologies do both current and lagged government R&D have a positive effect. However, the magnitude of the elasticities is quite small, suggesting that the effect was minimal. The sum of the elasticities for current and lagged government R&D for the supply technologies is only 0.265. To put this in perspective, consider that the average R&D expenditure per new patent is \$1 million. In contrast, federal R&D spending must increase by over \$100 million to induce a new patent!

IV. Conclusions

This paper uses patent data to study the impact of energy prices on energy-saving technology. It adds to the literature on induced innovation by looking not only at the effects of prices on technological change, but also at the effects of the usefulness of existing knowledge on technological change. The most significant result of this paper is the strong, positive impact of energy prices on new innovations. This suggests that environmental taxes and regulations not only reduce pollution by shifting behavior away from polluting activities, but also encourage the

³⁴ The change in focus on R&D suggests that a test for structural change would be appropriate. Unfortunately, the data do not offer enough degrees of freedom to perform such a test.

development of new technologies that make pollution control less costly in the long-run. My results also makes clear that simply relying on technological change as a panacea for environmental problems is not enough. There must be some mechanism in place that encourages innovation to occur.

My results also shed light on the determinants of technological change. In particular, I show that it is necessary to account for changes in the usefulness of the knowledge available to inventors in order to accurately estimate the effects of induced innovation. The *supply of existing ideas*, as well as the demand for new ideas, plays an important role in shaping the direction of innovation. Technologies for which little chance of successful innovation was possible did not experience significant shocks to innovation when energy prices were higher. Furthermore, I demonstrate that patent citations can be used as a measure of the supply of knowledge available to inventors when they undergo research. Finally, the productivity estimates obtained from the citation data exhibit a downward trend, suggesting diminishing returns to energy research over time.

REFERENCES

- Ahmad, Syd (1966), "On the Theory of Induced Innovation," *Economic Journal*, 76, 344-357.
- Binswanger, Hans P. (1974), "A Microeconomic Approach to Innovation," *Economic Journal*, 84, 940-958.
- Binswanger, Hans P. (1978a), "The Microeconomics of Induced Technical Change," in *Induced Innovation: Technology, Institutions, and Development*, ed. by H. Binswanger and Vernon Ruttan, John Hopkins University Press, Baltimore, MD.
- Binswanger, Hans P. (1978b), "Issues in Modeling Induced Technical Change," in *Induced Innovation: Technology, Institutions, and Development*, ed. by H. Binswanger and Vernon Ruttan, John Hopkins University Press, Baltimore, MD.
- 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*, MIT Press, Cambridge, MA.
- Cohen, Linda R. and Roger G. Noll (1991), *The Technology Pork Barrel*, The Brookings Institution, Washington, D.C.
- Evenson, Robert E. (1991), "Patent Data by Industry," Cowles Discussion Paper no. 620, Yale University.
- 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.
- Greene, William H. (1993), *Econometric Analysis*, Macmillan Publishing Co., New York.
- Griliches, Zvi (1990), "Patent statistics as economic indicators: a survey," *Journal of Economic Literature*, 28, 1661-1707.
- Griliches, Zvi (1989), "Patents: Recent trends and puzzles," *Brookings Papers on Economic Activity: Microeconomics*, 291-330.
- Hicks, John R. (1932), *The Theory of Wages*.
- 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, 577-598.

- Jaffe, Adam B. and Karen Palmer (1997), "Environmental Regulation and Innovation: A Panel Data Study," *Review of Economics and Statistics*, 79, 610-619.
- Jaffe, Adam B. and Manuel 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, 12671-12677.
- Kamien, Morton I. and Nancy L. Schwartz (1968), "Optimal Induced Technical Change," *Econometrica*, 36, 1-17.
- Kortum, Samuel and Josh Lerner (1998), "Stronger Protection or Technological Revolution: What is Behind the Recent Surge in Patenting?" *Carnegie-Rochester-Conference-Series-on-Public-Policy*, 48, 247-304.
- Lanjouw, Jean O. and Ashoka Mody (1996), "Innovation and the International Diffusion of Environmentally Responsive Technology," *Research Policy*, 25, pp. 549-571.
- 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.
- Mowrey, David C. and Nathan Rosenberg (1979), "The Influence of Market Demand Upon Innovation: A Critical Review of Some Recent Empirical Studies," *Research Policy*, 8, 103-153.
- Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins (1998), "The Induced Innovation Hypothesis and Energy-Saving Technological Change," NBER working paper #6437.
- Nordhaus, William (1997), "Modeling Induced Innovation in Climate-Change Policy," paper presented at the 1997 NBER Summer Institute on the Environment.
- Nordhaus, William (1973), "Some Skeptical Thoughts on the Theory of Induced Innovation," *Quarterly Journal of Economics*, 87, 208-219.
- Office of Technology Assessment (1991), *Energy Technology Choices: Shaping Our Future*, OTA-E-493, Washington, DC, U.S. Government Printing Office, July.
- Popp, David (1998), "The Productivity of Energy Research," presented at the 1998 NBER Summer Institute on Productivity.
- Public Citizen (1992), "Renewable Energy Research and Development: An Alternative Budget Proposal for FY 1993-1995."
- Rosenberg, Nathan (1984), *Exploring the Black Box: Technology, Economics, and History*, Cambridge University Press, Cambridge, U.K.

Rosenberg, Nathan (1982), *Inside the Black Box: Technology and Economics*, Cambridge University Press, Cambridge, U.K.

Scherer, F.M. (1965), "Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions," *American Economic Review*, 55, pp. 1097-1125.

Schmookler, Jacob (1966), *Invention and Economic Growth*, Harvard University Press, Cambridge, MA.

U.S. Department of Energy (1987), *Energy Conservation: Multi-Year Plan FY 1989-93*.

Table 1 – Technology Groups Included in the Energy Data Set**Supply Technologies:**

Coal liquefaction – producing liquid fuels from coal
Coal gasification – producing gaseous fuels from coal
Solar energy
Batteries for storing solar energy
Fuel Cells
Wind energy
Geothermal energy
Using waste as fuel
Using waste gases as fuel
Ocean Thermal Energy Conversion (OTEC)
Renewable energy – general

Demand Technologies

Recovery of waste heat for use as energy
Heat exchange – refrigeration
Heat exchange – general
Heat pumps
Stirling engine
Continuous casting processing of metal
Carbothermic processes for the manufacture of aluminum
Electrolysis for the manufacture of aluminum
Use of black liquor in paper manufacturing
Insulated windows
Compact fluorescent lightbulbs*

* Note: no classification could be found for these. Patents for compact fluorescent lightbulbs were found by a keyword search on the MicroPatent database. Because this is a relatively new technology, it was not necessary to extend the search before 1975.

Table 2 – Summary Patent Data
Privately Held U.S. Patents
 Sorted By Year of Application

Year	Coal Liquefaction	Coal Gasification	Solar Energy	Solar Batteries	Fuel Cells	Wind Power	Geothermal Energy	Waste as Fuel	Waste Gas as Fuel	OTEC	Renewable Energy – General	Waste Heat	Heat Exchange – Refrigeration	Heat Exchange – General	Heat Pumps	Stirling Engine	Continuous Casting	Aluminum – Carbothermic	Aluminum – Electrolysis	Black Liquor	Insulated Windows	CFL
1970	42	14	6	18	43	2	3	63	4	0	1	17	11	425	0	13	84	2	8	4	5	0
1971	37	24	5	17	46	8	4	53	5	1	2	18	13	423	2	13	115	0	9	2	3	0
1972	27	16	10	13	33	7	5	52	12	0	0	21	11	340	8	12	67	1	10	3	8	0
1973	28	20	36	23	28	4	3	52	14	0	0	12	18	346	4	9	63	2	9	2	7	0
1974	51	38	104	27	26	20	22	49	7	2	1	28	15	382	7	17	48	2	4	2	12	2
1975	45	31	218	63	38	37	17	29	13	3	7	26	20	418	8	11	46	2	2	3	17	0
1976	107	42	321	89	32	24	15	32	11	9	4	34	18	450	20	17	43	2	11	0	11	0
1977	88	45	367	117	52	29	14	34	14	3	7	29	19	505	17	11	37	3	10	3	9	1
1978	114	53	333	142	42	39	15	41	6	7	5	16	22	479	32	12	40	0	8	3	9	7
1979	77	32	295	119	40	33	11	40	9	6	5	27	31	462	24	11	45	1	11	1	3	10
1980	97	38	278	112	54	34	15	50	4	7	3	25	13	443	21	18	44	4	15	3	13	3
1981	100	27	208	119	54	37	14	44	3	3	1	23	22	382	30	21	43	1	12	1	8	4
1982	82	25	151	93	74	28	12	58	7	2	1	31	11	377	18	30	49	1	16	2	12	4
1983	74	22	102	74	47	23	12	50	4	2	1	22	10	317	11	21	61	2	14	1	7	6
1984	70	15	104	86	39	13	2	44	4	1	1	24	9	338	8	19	62	2	12	2	6	9
1985	34	18	85	80	54	9	5	46	6	0	0	17	15	286	14	13	46	0	10	3	4	2
1986	20	10	42	73	72	11	4	61	3	1	0	13	23	323	15	13	80	6	5	1	12	1
1987	12	16	35	54	65	9	6	83	8	2	0	13	17	297	11	19	39	0	6	0	11	5
1988	14	10	44	63	54	6	4	69	2	0	2	26	24	315	5	10	58	0	5	4	9	4
1989	22	14	33	42	51	6	6	84	10	0	0	24	22	311	14	12	38	0	5	3	20	6
1990	16	9	26	41	60	6	4	102	5	1	2	24	27	337	18	18	33	0	5	2	18	3
1991	10	4	32	48	49	15	4	98	12	0	0	19	35	391	22	11	38	0	6	2	7	5
1992	12	5	27	53	61	13	11	93	9	2	1	25	28	428	14	12	33	0	1	6	16	12
1993	8	2	23	20	58	10	2	60	17	0	0	22	13	350	17	5	31	2	3	2	2	7

Table shows the number of successful patent applications in each technology field by U.S. inventors.

As discussed in the text, data after 1985 have been scaled up to include applications not yet acted upon by the U.S. Patent Office.

Table 3 – Correlations Between Energy Prices and Patent Counts
U.S. patents sorted by year of application

Technology	<i>correlation with:</i>		
	current prices	prices lagged 1 year	prices lagged 2 years
Coal Liquefaction	0.424	0.251	0.034
Coal Gasification	0.059	-0.179	-0.299
Solar Energy	0.325	0.100	-0.148
Solar Batteries	0.675	0.548	0.331
Fuel Cells	0.517	0.611	0.645
Wind Power	0.477	0.223	-0.093
Geothermal Energy	0.246	-0.065	-0.323
Waste as Fuel	-0.073	0.028	0.162
Waste Gas as Fuel	-0.413	-0.544	-0.567
Ocean Thermal Energy Conversion	0.233	0.008	-0.153
Renewable Energy – General	-0.036	-0.189	-0.329
Waste Heat	0.283	0.055	-0.151
Heat Exchange – Refrigeration	-0.030	-0.110	-0.177
Heat Exchange – General	-0.175	-0.297	-0.413
Heat Pumps	0.544	0.373	0.120
Stirling Engine	0.662	0.597	0.463
Continuous Casting	-0.442	-0.209	0.142
Aluminum -- Carbothermic	0.107	0.145	0.010
Aluminum – Electrolysis	0.525	0.424	0.340
Black Liquor	-0.287	-0.181	-0.199
Insulated Windows	0.093	-0.047	-0.184
Compact Fluorescent Lightbulbs	0.548	0.394	0.412

Table presents correlations between energy prices and the number of patent applications per year. Only privately held patents by American inventors are included.

Table 4 – Induced Innovation Regression Results

Dependent variable: percentage of total domestic patent applications in each technology group

Independent Variables	All Technologies	Supply Technologies	Demand Technologies
Energy Prices	0.125 (2.862)	1.080 (5.414)	0.801 (6.000)
Lagged Energy Prices	0.462 (10.795)	-0.887 (-4.488)	0.268 (1.998)
Marginal Productivity of R&D	0.259 (23.668)	0.115 (2.978)	1.865 (40.352)
Time	-0.152 (-74.959)	-0.425 (-52.123)	-0.306 (-33.571)
Government R&D	0.010 (7.279)	0.018 (3.033)	0.039 (8.047)
Lagged Government R&D	-0.016 (-11.551)	0.247 (4.639)	-0.041 (-8.255)
Adjusted R-square	0.775	0.950	0.782
number of technology groups:	17	9	8

*t-stats below estimates**T = 20*

Table shows the induced innovation regression results. Lagged party of the President and lagged government R&D used as instruments for government R&D.

Table 5 – Induced Innovation Regression Results – No Productivity Estimates
 Dependent variable: percentage of total domestic patent applications in each technology group

A. No control for productivity of R&D (no time trend or productivity estimates)

Independent Variables	All Technologies	Supply Technologies	Demand Technologies
Energy Prices	1.289 (13.687)	0.903 (2.251)	0.008 (0.032)
Lagged Energy Prices	-0.769 (-8.364)	-0.328 (-0.859)	0.104 (0.412)
Government R&D	-0.009 (-3.252)	0.020 (2.119)	0.034 (3.491)
Lagged Government R&D	0.007 (2.635)	0.007 (0.888)	-0.015 (-1.693)
Adjusted R-square	0.866	0.840	0.944
number of technology groups:	17	9	8
F-test for H_0 :			
Productivity of R&D = time = 0	3149.192	1358.653	1005.002

B. Time trend used instead of productivity parameters

Independent Variables	All Technologies	Supply Technologies	Demand Technologies
Energy Prices	1.896 (40.972)	1.629 (5.177)	-0.267 (-1.549)
Lagged Energy Prices	-0.459 (-10.051)	0.355 (1.135)	-2.329 (-13.333)
Time	0.556 (9.574)	-0.749 (-2.371)	-0.659 (-2.856)
Government R&D	-0.032 (-20.723)	0.012 (1.651)	0.281 (5.402)
Lagged Government R&D	-0.005 (-3.694)	-0.034 (-5.126)	0.112 (22.621)
Adjusted R-square	-0.683	0.695	0.232
number of technology groups:	17	9	8

t-stats below estimates

$T = 20$

Table shows the results of the induced innovation regressions when productivity estimates are removed. Part A includes no control for the productivity of R&D. Part B uses a time trend in place of the productivity estimates for all technology groups. Note that the price elasticities are lower and that the fit of the regressions is worse than when the productivity estimates are included.

Figure 1 – Successful Patent Applications by U.S. Inventors






Figure shows the number of successful domestic patent applications in five technology groups. Energy prices are the cost per million Btu of energy consumption. All data have been normalized so that 1981 = 100.

Figure 2 – Probability of Citation



The figure presents the probability that patents granted in year x will be cited by patents applied for in subsequent years. Each line represents the patents granted in a different year. The x -axis is the number of years since the patent was granted.


Figure 3 – Productivity Estimates



Figures show the productivity estimates for each technology group, with 1970 normalized to 1 in each case. Note that, for most technologies, there is a declining trend to the estimates, suggesting diminishing returns to research over time.

Figure is continued on the next page.

Figure 3 – Productivity Estimates (continued)



<i>Other Results</i>					
citing practices:			group dummies:		
	estimate:	standard error:		estimate:	standard error:
1974-1975	1.000	N/A	coal liqefaction:	4.621	1.144
1976-1977	0.906	0.056	coal gasification:	4.110	1.681
1978-1979	0.839	0.073	solar energy:	1.208	0.576
1980-1981	0.777	0.090	solar batteries:	2.670	0.764
1982-1983	0.846	0.124	fuel cells:	1.374	0.426
1984-1985	0.915	0.165	wind:	12.468	4.627
1986-1987	0.934	0.200	waste fuel:	1.557	0.609
1988-1989	0.968	0.240	waste heat:	1.538	1.180
1990-1991	0.993	0.279	heat exchange:	0.172	0.065
decay:	0.392	0.018	heat pumps:	10.362	3.443
diffusion:	0.00251	0.00061	Stirling engines:	5.926	1.772
adjusted R-square:	0.778		continuous casting:	1.000	N/A

Figures show the productivity estimates for each technology group, with 1970 normalized to 1 in each case. Note that, for most technologies, there is a declining trend to the estimates, suggesting diminishing returns to research over time.