NBER WORKING PAPER SERIES

INDUCED INNOVATION AND ENERGY PRICES

David Popp

Working Paper 8284
http://www.nber.org/papers/w8284

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2001

I would like to thank William Nordhaus, Joel Waldfogel, Robert Evenson, Adam Jaffe, Stu Bretschneider, Ariel Pakes, Dietrich Earnhart, Tom Weiss, and three anonymous referees for helpful comments. Despite their useful suggestions, however, all mistakes are my responsibility. In addition, I also thank Adam Jaffe for providing the programming code for the patent citation regressions. Finally, I thank Robert Evenson, Jon Putnam, and Samuel Kortum for their help in obtaining patent data. Partial financial support for this article comes from DOE grant 593A3140217. The views expressed herein are those of the author and not necessarily those of the National Bureau of Economic Research.

Induced Innovation and Energy Prices
David Popp
NBER Working Paper No. 8284
May 2001
JEL No. O31, Q40, Q42

ABSTRACT

I use U.S. patent data from 1970 to 1994 to estimate the effect of energy prices on energy-efficient innovations. Using patent citations to construct a measure of the usefulness of the existing base of scientific knowledge, I consider the effect of 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. I find that both energy prices and the quality of existing knowledge have strongly significant positive effects on innovation. Furthermore, I show that omitting the quality of knowledge adversely affects the estimation results.

David Popp
Department of Public Administration
Center for Environmental Policy Administration
Center for Technology and Information Policy
The Maxwell School
Syracuse University
400 Eggers Hall
Syracuse, NY 13244-1090
and NBER
Tel: 315-443-2482
Fax: 315-443-1075
Email: dcpopp@maxwell.syr.edu
Web: http://faculty.maxwell.syr.edu/dcpopp/index.html

In recent years, environmental economists have paid increased attention to the possibilities that price-induced technological change might introduce for environmental policy. Many environmental problems, such as global warming, are long-term problems for which technological progress may play a crucial ameliorating role. Economists often cite induced technological change as an advantage to using market-based environmental policies for dealing with these problems. Unfortunately, there is little empirical evidence about the policy-induced development of environmentally friendly technology.[1] As a result, most environmental policy models treat technological change as exogenous and therefore cannot examine potentially important links between policy and technological change.[2] Furthermore, models that do incorporate endogenous technological change have been hampered by a lack of empirical evidence on the links between environmental policy and innovation.[3]

In order to address the shortcomings in our understanding of endogenous technological change, I use U.S. patent data from 1970 to 1994 to study the effect of energy prices on energy-efficient innovations. Although my main focus is the influence of energy prices on innovative activity related to energy efficiency, I also hope to shed light on the more general question, What factors influence inventive activity? In 1932, J. R. Hicks introduced the theory of induced innovation, which states that changes in relative factor prices should lead to innovations that

---

[1] The existing literature on induced environmental innovation includes Jean O. Lanjouw and Ashoka Mody (1996), Adam B. Jaffe and Karen Palmer (1997), and Richard G. Newell *et al.* (1999). However, only Jaffe and Palmer include elasticities that could be incorporated into large-scale models of the environment.

[2] Two recent exceptions are Lawrence H. Goulder and Stephen H. Schneider (1999) and William Nordhaus (1999). As evidence that there is still much work to be done in correctly modeling induced technological change, the authors of the two articles 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."

reduce the need for the relatively expensive factor. Interest in the microeconomics of induced innovation increased after articles by Syed Ahmad (1966), Morton Kamien and Nancy Schwartz (1968), and Hans Binswanger (1974, 1978a, 1978b) were published, but empirical testing of the induced innovation hypothesis was hampered by a lack of available data. My use of the U.S. patent data overcomes this deficiency.

Because the induced innovation literature treats the existing stock of knowledge on which inventors can build as exogenous, it ignores the determinants of that base of knowledge and therefore cannot fully 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 research and development (R&D) may make future R&D more difficult. By contrast, *technology-push* theories of R&D emphasize the importance of technological opportunity to innovation, and thus they are able to capture the links between current and future research.[4] In this paper, I make use of patent citation data to measure technological opportunity in order to combine information on both market demand and technological opportunity in an empirical study of the determinants of innovation. In so doing, I hope to show that including information on both is crucial for an accurate estimation of the factors that drive innovation.

In the first section, I describe the construction of the energy patent data set and discuss how it will be used. In addition to data on energy patents and prices, a measure of the usefulness of the existing stock of scientific knowledge must also be included. In the second section, using patent citations, I construct *productivity estimates* that capture changes over time in the

[4] For examples, see F.M. Scherer (1965), Jacob Schmookler (1966), Nathan Rosenberg (1982, 1984), and David C. Mowrey and Rosenberg (1979).

usefulness of energy patents for future inventors. I combine these productivity estimates with patent data to construct stocks of knowledge available to inventors in each year. These stocks are used in the third section, along with the rest of the data in the first section, 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 stock of knowledge available to inventors have strongly significant positive effects on innovation. Furthermore, I show that both not adjusting patents in the stocks for quality and omitting the stocks completely lead 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 and 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 at that time. As a result, most existing tests of the induced innovation hypothesis have been indirect tests focusing on the results of innovation rather than on the innovation process itself. In creating a data set of energy patents in the United States from 1970 to 1994, made possible by the recent computerization of patent data, I have been able to develop a direct test of the induced innovation hypothesis.

A. The Energy Patent Data Set

All patents granted in the United States are given a U.S. classification number. There are currently over 300 main classification groups and over 50,000 subclassifications. In order to identify subclassifications pertaining to energy efficiency, I used resources from the Department of Energy and from the academic sciences to identify several areas of research in the energy field. Descriptions of these technologies were then matched with U.S. patent subclassifications, and those technologies for which no clear subclassification existed were eliminated. The resulting set of subclassifications was then sorted into 11 distinct technology groups, including 6 groups pertaining to energy supply, such as solar energy, and 5 groups relating to energy demand, such as methods of reusing industrial waste heat. The first column of Table 1 lists the 11 energy supply technology groups.[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 11 technology groups that were granted in the United States between 1950 and 1994.[6] For the purposes of this study, only patents granted to Americans were included, since foreign inventors are likely to be influenced by conditions 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 were not included.

Patents were sorted by the year of application,[7] and since information on patents is not made public in the United States until the patent is granted, only *successful* patent applications are included in the data set. Several researchers have found that grouping patents by the date of application is a good indicator of R&D activity (for example, see Zvi Griliches 1990), but since a

---

[5] An earlier version of this paper included 22 different technology groups. To construct the knowledge stocks in section II, it was necessary to remove groups that did not have a consistent flow of patents over the past 50 years. Interested readers may download a more thorough description of the technologies chosen, as well as data on these additional technology groups, at http://faculty.maxwell.syr.edu/dcpopp/index.html.

[6] The MicroPatent database contains every U.S. patent issued from 1975 to 1994. To consider data from the first energy crisis of 1973, as well as to have enough patents to construct stocks of existing knowledge, additional patent data were obtained from the Classification and Search Support System (CASSIS) available from the U.S. Patent and Trademark. Unfortunately, some of the data found in the MicroPatent database, such as patent citations, are not available from these other sources.

[7] Note that the patent application date was 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 random sample of patents in the data set indicates that for about 20 percent of the patents, the original application date would be earlier than the application date listed 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. As I nonetheless find significant coefficients in the regressions, the effect of this error appears to be small.

patent application is made public only when a patent is granted, the data have been scaled up to account for patents applied for but not yet granted. First, the distribution of the lag between application and grant for all energy patents in the data set was found, and using this distribution, the percentage of patents still waiting to be granted in 1994 was calculated 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 it has not yet been possible to cite recently granted patents, only patent applications through 1990 are used in the subsequent regression analysis.[8]

Table 1 shows the annual count of successful patent applications in each of the technology groups from 1970 to 1993.[9] Data for the years after 1985 have been scaled to include applications not yet acted upon, as described above. Figure 1 illustrates trends in energy prices and data for five of the technology groups, with 1981 normalized to 100.[10] The cost of energy is shown in constant 1987 dollars per million British thermal units (Btus). For most of the technology groups, there was a jump in patent applications during the energy crises of the 1970's, which suggests that energy prices do play an important role in inducing energy-efficient technological change.[11] This relationship is further evident in Table 2, 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.

[8] The scaling should have little impact on the final results. Given the distribution of lags between application and grant in the data, 99 percent of the patents applied for in 1990 have already been granted. Thus, no figure used in the regression has been increased by more than 1 percent.

[9] Although only patent applications through 1990 are used in the subsequent regressions, the figures include patent applications through 1993 for informational purposes.

[10] Additional figures are available upon request.

[11] 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.

In addition to illustrating the strong effect that energy prices have on patenting activity, Figure 1 also highlights the importance of considering the possibility of diminishing returns to R&D. Note that energy prices do not peak until 1981 but patenting activity in most of these technology groups reached 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. The fact that patenting activity dropped before prices suggests that the possibility of diminishing returns to research should be explored.[12]

B. Other Data

In addition to data on energy patents, other data are also used. Data on energy prices are taken from the *State Energy Price and Expenditure Report* published by the Energy Information Administration, and are available since 1970. Prices are in constant 1987 dollars per million Btus deflated by a gross-national-product 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, there is also an index of total energy prices. Data for government R&D are taken from the interest group Public Citizen (1992), which collected data on federal R&D expenditures by energy technology from annual U.S. government budget publications.

I also include data for several technology-specific variables, which capture the effects of characteristics unique to the individual technology groups and help to explain some of the individual trends mentioned in footnote 11. 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.

---

[12] One might argue that another possible explanation for the earlier decline in energy patents is that political support for energy R&D may have changed. However, the drop in patenting occurred during the years of the Carter administration, when support for energy research was at its highest.

With regard to the use of waste products as energy, the price utilities pay to purchase waste products for fuel is included in the data set. This figure captures the increased supply of waste that became available as fuel owing to concerns about declining landfill space during the 1980's. Finally, for the fuel cells group, 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 is 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 overestimated future electricity demand in the 1970's, the electric utility industry built too many new large power plants, and fuel cells offered a more flexible alternative when increased power-generating capacity was needed in the 1980's (Office of Technology Assessment 1991). As a consequence, we would expect a positive correlation between excess capacity and fuel cell patents.

C. Modeling

I use the data just described to test the induced innovation hypothesis for energy research. Because much of the R&D process is poorly understood, specifying a structural model of the determinants of energy patents is difficult. Instead, I use a simple log-log regression of successful patent applications on energy prices, a knowledge stock, and the other variables discussed in the previous subsection. Such a specification allows the resulting coefficients to be interpreted as elasticities, which should be of use to modelers of environmental problems such as global warming. Since lagged values of the prices are important to the expectations of price at any given time, and since the price data begins just three years before the beginning of the energy crisis in 1973, I use a distributed lag model to model the effect of prices and government R&D.

$EPAT_{i,t}$ represents the number of successful nongovernment U.S. patent applications for technology field *i* in year *t,* and $TOTPAT_t$ represents the total number of successful nongovernment U.S. patent applications in the same year. $P_{E,t}$ is the price of energy in that year. The variable $K_{i,t\text{-}1}$ represents the stock of knowledge that had accumulated by the previous year, and should be thought of as the knowledge available to the researcher at time *t*. Values for this stock are calculated in the next section. **Z** is a vector of the other independent variables described in the previous subsection, such as R&D spending by the U.S. Department of Energy. For any energy-saving technology, *i*, I estimate the following model:

(1) $$\log\left(\frac{EPAT_{i,t}}{TOTPAT_t}\right) = \varphi_i + \gamma(1-\lambda)\log P^*_{E,t} + \theta \log K_{i,t-1} + \boldsymbol{\eta}(1-\lambda)\log \mathbf{Z}^*_{\mathbf{i,t}} + \lambda^t \mu^0 + \varepsilon_{it}\,,$$

$i = 1,\ldots,11;\ \ t = 1,\ldots,20$

where:

(2) $$P^*_{E,t} = P_{E,t} + \lambda P_{E,t-1} + \lambda^2 P_{E,t-2} \cdots + \lambda^{t-1} P_{E,1}, \text{ and}$$

(3) $$\mathbf{Z}^*_{\mathbf{i,t}} = \mathbf{Z}_{\mathbf{i,t}} + \lambda \mathbf{Z}_{\mathbf{i,t-1}} + \lambda^2 \mathbf{Z}_{\mathbf{i,t-2}} \cdots + \lambda^{t-1} \mathbf{Z}_{\mathbf{i,1}}\,.$$

This model is consistent with an adaptive expectations model of prices, in which expected future prices depend on a weighted average of past prices. λ, the adjustment coefficient, represents the weights placed on past observations. μ^0 represents the truncation remainder, since an infinite series of past independent variables is not possible. In this model, $\gamma(1\text{-}\lambda)$ represents the short run price elasticity of energy innovation, and γ represents the long run elasticity. Finally, because decisions about government funding of energy R&D are correlated with energy prices, instrumental variables are used for federal energy R&D spending, including a variable for lagged federal energy R&D and a dummy variable representing the lagged political party of the president.[13]

Note that the model uses the *percentage of all successful domestic patent applications per year in each technology field* as the dependent variable. Using the percentage of applications in each field, rather than a raw count of applications, accounts for growth in the economy and exogenous changes in patenting behavior.[14] Changes that affect *all* patent classifications would lead to a change in both the total number of patent applications and the number of energy patent applications in a given year. For example, a change in patent law that increases the propensity to patent by 5 percent would increase both the numerator and the denominator of the dependent variable by 5 percent.

Using patent data as a measure of research output poses several complications. The first is that the propensity to apply for patents 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 the latter industries, the cost of revealing an idea to competitors is often greater than what is gained by having patent protection. As a result, the correlation between R&D and patents varies across industries.[15] However, I am more concerned with the time series aspects of the data, and in that case, 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 when there is an increase in energy prices. As long as the tendency to

---

[13] 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.

[14] During the years studied, total domestic patent applications numbered 54,894 in 1971, fell to a low of 31,548 in 1985, and rose to 59,032 by 1990.

[15] Richard C. Levin *et al.* (1987) discuss the variation in patenting behavior across industries.

patent remains the same in each field across time, variations in the tendency to patent across different fields do not pose a problem.

Unfortunately, a second difficulty with using patent data is that it is not clear that the tendency to patent is the same across 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 Robert E. Evenson (1991), consider the falling ratio to be evidence of diminishing returns to R&D.[16] Similarly, Samuel Kortum and Josh 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 Zvi Griliches (1989), claim that research opportunities have not declined. Griliches argues that the fall in the patent-to-R&D ratio is a result of changes in the willingness of inventors to patent new inventions. Exogenous factors that caused a fall in the willingness to patent – for example, changes in the 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 in examining the falling patent-to-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 the R&D. The model to be estimated includes two features to help identify and control for the cause. First, by using the percentage of all successful domestic patent applications per year in each technology field as the dependent variable, I eliminate changes that affect *all* patent classifications. For example, the Bayh-Dole Act of 1980

---

[16] In this paper, diminishing returns to research refers to the expected return on the *inputs* to the research process, not the returns to the output. 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.

allowed universities to patent inventions, which resulted in a large increase in the number of patents granted in the U.S. Since such a change increases both the numerator and the denominator of the dependent variable *by x* percent, using the percentage of patent applications removes exogenous changes in the propensity to patent cancel each other out and are removed from the data.

The second step in identifying changes in patenting behavior over time is to include a measure of the stock of knowledge available to inventors at time *t*. I construct such a stock in the next section. Using patent citations to measure the usefulness of existing patents to future researchers, I find evidence that the productivity of existing knowledge to future researchers falls over time, suggesting that there are diminishing returns to the output of research 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 stock of knowledge suggests that diminishing returns to research explain at least a part of the decline in the patent-to-R&D ratio.

II. Patent Citations and the Existing Stock of Knowledge

In this section, I use data on patents granted in each technology group since 1950 to construct stocks of knowledge available to inventors in each year from 1970 to 1991. First, I use patent citation data to ascertain the usefulness of patents granted in each group in each year to future inventors. I then combine these estimates of the usefulness of patents with patent counts from 1950-1990 to construct three different measures of the stock of knowledge.

A. Citations and the Productivity of Knowledge

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, and 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.[17] 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.

For this paper, I use patent citations as evidence of the existing state of technology when the invention was completed. The assumption is that the citations indicate a flow of knowledge. Thus, 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.[18] We would expect that the marginal productivity of research would be greater in the years immediately after useful patents were revealed.

In making a formal analysis of patent citations, a simple count of subsequent citations is not enough, for 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 for such factors.

---

[17] "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.

[18] Jaffe, *et al.* (1998) examined the relationships between knowledge flows and patent citations. Their research included interviews with scientists, R&D directors, and patent attorneys. They found that, at the level of individual patents, not all citations are indicative of knowledge flows, as other concerns, such as strategically including irrelevant patents to satisfy the patent examiner, affected the citation process. However, on more aggregate levels, such as the patents for an organization or a firm, they found that patent citations do seem indicative of knowledge flows. Similarly, Lanjouw and Schankerman (1999) find that forward citations (citations made by future patents to an existing patent) are one of the least noisy indicators of the quality of the existing patent.

For each technology group, potentially cited patents are sorted by the year in which the patent was granted and are denoted *CTD,* for "year cited." Since a patent application is not made public in the U.S. unless the patent is granted, 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 described in the preceding section.

In principle, regression analysis could be used on each individual patent in the data set, but the dependent variable – the actual number of citations – is zero for most patents, as most patents are never cited. As a result, the data are sorted by cited and citing years into groups of patents that could potentially cite each other. Separate groups are constructed for each technology group, i, and only citations made by other patents in the technology group are considered.[19] A total of 6,417 i,CTD,CTG groups resulted. As an example of such a group, one might be made up of citations to all solar energy patents granted in 1975 made by solar energy patents applied for in 1980.[20] Denoting citations in each group as $c_{i,CTD,CTG}$, the number of potentially cited patents applied for in year CTD as $n_{i,CTD}$, and the number of potentially citing patents granted in year CTG as $n_{i,CTG}$, the probability of citation, p, for patents within each group

---

[19] As discussed in Popp (1998), similar results are obtained when citations to all patents are considered. Including citations to all patents would allow for the possibility of spillovers of knowledge, but because spillovers do not necessarily impact R&D in the energy sector right away, to include citations to all patents would complicate the induced innovation regression.

[20] Although the data set includes patents granted since 1950, citations are available only for patents granted after 1975. Because of the lag between the application and the granting of a patent, the average being two years, I am able to consider citations made by all patents applied for in 1974 or later.

is:

(4) $$p_{i,CTD,CTG} = \frac{c_{i,CTD,CTG}}{(n_{i,CTD})(n_{i,CTG})},$$

To control for factors that affect the likelihood of citation, I used a model first used by Ricardo J. Caballero and Adam B. Jaffe (1993) and Jaffe and Manuel Trajtenberg (1996) to estimate the probability that a patent would be cited by subsequent patents. The model uses an exponential distribution to model flows of knowledge. In this model, the probability of citation is written as:

(5) $$p(i,CTG,CTD) = \alpha(i,CTG,CTD)\exp[-\beta_1(CTG\text{-}CTD)][1\text{-}\exp(-\beta_2(CTG\text{-}CTD))].$$

β_1 represents the rate of decay of knowledge as it becomes obsolete, and β_2 stands for the rate at which newly produced knowledge, as represented by a newly patented innovation, diffuses through society. Parameters capturing attributes of the citing or cited patents that may influence the probability of citation, $\alpha(i,CTG,CTD)$,[21] include:

- the usefulness of the knowledge represented in the patent being cited ($\alpha_{i,CTD}$),
- the frequency with which patents applied for in the citing year cite earlier patents (α_{CTG}),[22] and
- the frequency of citations within each technology group (α_i).

The first two α parameters are based on Jaffe's earlier work; α_i is added for this study to

[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 nonlinearly.

[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*.

capture the effect of the size of the technology group. About half of all patent citations are to patents in the same classification (Jaffe *et al.* 1993), but the technology groups I am using 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, which means that citations to other patents in the group are less likely in those groups.

The first parameter, $\alpha_{i,CTD}$, the *productivity parameter,* is the value of interest here, for it tells us the likelihood that patents from the 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 $\alpha_{i,CTD}$ indicate that the patents in question are more likely to be cited by subsequent patents, which implies that the knowledge embodied in those patents is particularly useful. If that stock of knowledge is particularly useful to inventors, research should be more productive.

Figure 2 provides examples of the trends in citation data for solar energy and fuel cells. The graphs plot the probability that patents granted in a particular year will be cited by other patents applied for in the same technology group *X* years after the year the patent is granted. Each line represents 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 one year being cited by a patent granted *X* years later. 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. As a result, the productivity parameter, $\alpha_{i,CTD}$, can be visualized as the y-intercept in the figures. Patents from years that are cited more often, such as the 1975 solar energy patents, have higher *y*-intercepts. Also note that the probability of citation falls over time, which suggests that the decay of knowledge is more influential than the diffusion of knowledge in determining the probability of citation.[23]

To estimate the values of the productivity parameter, $\alpha_{i,CTD}$, insert the variables defined above and add an error term, $\varepsilon_{i,CTD,CTG}$, to equation (5), which yields:

$$p_{i,CTD,CTG} = \alpha_i\, \alpha_{i,CTD}\alpha_{CTG}\exp[-\beta_1(CTG\text{-}CTD)]\{1\text{-}\exp[-\beta_2(CTG\text{-}CTD)]\} + \varepsilon_{i,CTD,CTG}. \tag{6}$$

This equation is estimated using nonlinear least squares, using all patents granted from 1950 to 1989 as the cited years, and using all patents applied for from 1974 to 1991 as the citing years.[24] To identify the parameters, it is necessary to normalize one of each of the αs to be 1—for cited years, α_{1970} is normalized to 1, for citing years, $\alpha_{1974\text{-}75}$ is normalized to 1, and α_i is normalized to 1 for continuous casting patents. Finally, since this is grouped data, the observations are weighted by $\sqrt{(n_{i,CTD})(n_{i,CTG})}$ to avoid problems with heteroskedasticity (Greene 1993).

Figure 3 displays the results of my estimation.[25] A plot of the productivity parameter is presented for each of the eleven technologies. Numerical results for the remaining parameters are given in Table 3. Readers interested in a more detailed discussion of these other parameters are referred to David Popp (1998). To interpret the plots of the productivity parameter, recall that the productivity parameters in 1970 are normalized to 1. Therefore, estimates greater than 1 for a

[23] In both halves of Figure 2, the probability of citation is highest in the year after the patent is granted, which is consistent with the data in Jaffe and Trajtenberg (1996). In their article, 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 α's for every possible year. Since the main parameters of interest are the productivity parameters, $\alpha_{i,CTD}$, separate coefficients are obtained for each of these. The parameters for institutional changes are grouped by two-year periods, so that α_1 represents citation practices in 1974-1975; α_2, citation practices in 1976-1977; etc.

[25] To highlight the variation during the energy crisis, the figure only presents the productivity parameters from 1970 to 1990. In general, the estimates are greater than one for the years prior to 1970. Complete estimation results are available at http://faculty.maxwell.syr.edu/dcpopp/index.html.

given year indicate that patents granted in that year are more likely to be cited by future patents than patents that were granted in 1970, and estimates less than 1 indicate that those patents are less likely to be cited. During the years of peak patenting activity, there is a downward trend in the productivity parameter 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.[26]

B. Constructing the Knowledge Stocks

Before proceeding with the induced innovation regressions, I use productivity estimates to construct a *stock of knowledge* for each technology group. For each technology group, I construct three separate stocks. In each case, the rates of decay and diffusion estimated above are used. The three stocks are:

1) A simple stock of previously granted U.S. patents:

$$K_{i,t} = \sum_{s=0}^{t} PAT_{i,t} \exp[-\beta_1 (t-s)]\{1-\exp[-\beta_2 (t-s)]\}$$

2) A stock of patents weighted by the productivity estimates:

$$K_{i,t} = \sum_{s=0}^{t} \alpha_{i,s} PAT_{i,t} \exp[-\beta_1 (t-s)]\{1-\exp[-\beta_2 (t-s)]\}$$

3) A stock of the productivity estimates without patents

$$K_{i,t} = \sum_{s=0}^{t} \alpha_{i,t} \exp[-\beta_1 (t-s)]\{1-\exp[-\beta_2 (t-s)]\}$$

---

[26] Note that for technologies with limited R&D data available, such as solar technology, the productivity estimates are positively correlated with patent to R&D ratios – high productivity estimates occur before years when R&D successfully led to many patents, and low productivity estimates proceed years in which R&D was less successful. This result suggests that the productivity estimates do pick up variation in the productivity of future R&D. Also, one referee noted that the productivity estimates could be affected by changes in the propensity to patent *within technology groups*. Although it is impossible to separately identify such an effect in this regression, the correlation with the patent to R&D ratio suggests that this is not a problem here. Since changes in the propensity to patent should affect patents of marginal quality, which would be cited less, these changes should cause the productivity

Figure 4 plots the resulting stocks from 1970-1991. Because the scale of each stock varies, in the figure I normalize the stocks so that 1970 = 1. First, note that the unweighted knowledge stocks increase over time for each of the 11 groups. In comparison, the two stocks constructed using the productivity estimates tend to fall over time, although there is more variation in stock 2, which combines patents and the productivity estimates. This result suggests that the quality of the knowledge represented by these patents falls over time, as the value of the patent stocks falls more quickly when weighted by citations.[27] Thus, the distinguishing feature between stock one and the others is that stock 1 has no control for the quality of patents.

Stocks 2 and 3 are distinguished as follows. Because it includes both the productivity estimates and patent counts, stock 2 places more weight on years with many low quality patents than does stock 3, whereas stock 3 places more weight on years with a few high quality patents.[28] Because patenting activity tends to increase in the technology groups, while the productivity estimates tend to fall, stock 3, which only uses the productivity estimates, shows even stronger diminishing returns to knowledge. In the results section, I show that accounting for the returns to patents from each year is important, and that using a stock of knowledge based on specification 3 greatly improves the quality of the induced innovation regression results.

---

estimates to be negatively correlated with the patent to R&D ratio.

[27] One caveat to interpreting the results of diminishing returns is necessary. This paper focuses on returns to a narrowly defined group of technologies, and 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.

[28] For example, the productivity estimate for solar energy in 1975 is 3.3, and in 1985 is 0.76. However, in 1975, only 68 solar energy patents were granted, versus 251 in 1985. Using only the productivity estimate, knowledge from 1975 is just over four times more valuable. Using the product of the productivity estimates and patent counts, is just 17% more valuable.

III. Results

Having used estimates of the productivity of knowledge to construct the knowledge stocks, I now move to estimating the induced innovation relationship in equation (1), which I repeat here:

$$(1) \qquad \log\left(\frac{EPAT_{i,t}}{TOTPAT_t}\right) = \varphi_i + \gamma(1-\lambda)\log P^*_{E,t} + \theta \log K_{i,t-1} + \boldsymbol{\eta}(1-\lambda)\log \mathbf{Z}^*_{\mathbf{i,t}} + \lambda^t \mu^0 + \varepsilon_{it}\,,$$

Lagged values of the knowledge stocks just obtained are used as measures of the knowledge available to inventors at time *t,* making it possible to control for supply-side factors that affect the level of innovation as well as for demand-side factors such as energy prices. For concise presentation of results, I pooled the technology groups to obtain single estimates for each parameter.[29]

To estimate the model, note that it is linear once λ is defined. Thus, I search over the range $0<\lambda<1$ to find the λ that best fits the data. Because the underlying model is a non-linear model with instrumental variables, I use generalized method of moments (GMM) to find λ. The weighting matrix in the GMM criteria allows for correlation of the error term across technology groups and corrects the results for autocorrelation.

Table 4 presents the results of equation (1) for each of the three knowledge stocks. The parameter results shown are the short run elasticities for each variable. *T*-statistics for each coefficient are included in parentheses. The results highlight the importance of controlling for

---

[29] One might expect the stocks of knowledge to be correlated with other independent variables, such as energy prices. To check for this correlation, I ran additional regressions using lagged values of the other independent variables and a time trend as instruments for the knowledge stocks. A Hausman tests fails to reject the null hypothesis that the knowledge stock is exogenous for all regressions except with the unweighted stock of knowledge. As such, in the results reported below, I use instrumental variables for the unweighted stock of knowledge, but not for the stocks that include the productivity estimates.

the quality of knowledge available to inventors. The first column presents results using an unweighted count of past patents as the knowledge stock. Although the signs of the coefficients are as expected, the value of λ, 0.933, is unrealistically high, as it suggests an average lag of 10 years, and a median lag of nearly 14 years. The implication of this result is that just half of the effect of the 1973 oil price shock on innovation would have passed by 1987!

Using a weighted stock of patents to represent knowledge provides a more realistic estimate of λ, but does not improve the estimate of the price elasticity of patents. Column 2 presents the results of this regression. The estimated short-run price elasticities are similar to the unweighted regression, although they are insignificant. More importantly, λ is now a more reasonable 0.781, suggesting a median lag of 3.57 years, and a mean lag of 2.81 years. However, the combination of the small short run elasticity and the shorter lag structure produces a long run elasticity of patents with respect to energy prices of just 0.056.

Finally, column 3 presents results using a stock of the productivity estimates only. Recall that the advantage of this stock of knowledge is that it places more weight on years with a few high quality patents than a year with many low quality patents. In essence, stock two places more importance on having lots of patents to build on, even if each individual patent is of little use, and specification three places more importance on having quality patents to build on, even if the number of patents available is low. Energy prices play a more important role in this specification. The short-run elasticity of energy patents with respect to price is 0.125, and the long run elasticity is 0.625. Indeed, if there were diminishing returns to R&D, including patents in the stock of knowledge would bias the price elasticity downward, since the returns to R&D would be falling while energy prices were rising, due to the increased number of patents occurring after higher energy prices. In addition, the estimate of the price elasticity is much more precise, and is now significant at the one percent level.[30]

The mean and median lags suggest that the reaction to higher energy prices is fairly quick. Using specification (3), the mean lag is 3.11 years and the median lag is 4.00 years. 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 inventor's interest to file patent applications early in the inventive process so that an early priority date is established. Thus, a quick reaction is not unexpected.

A. 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 quality of the stock of knowledge available to the inventor is also an important factor. The coefficients of the lagged knowledge stock are positive and significant in each regression. Furthermore, combining the estimated elasticities with the variation in the knowledge stocks reveals that the magnitude of the productivity effect is large. The average change in the weighted knowledge stocks from year to year is 29 percent.[31] As a result, changes in the stock would change patenting activity in an average year by 24.19 percent. Using the stock of productivity estimates only, the average change in the stock is 9.5 percent, which leads to a 4.17 percent change in patenting activity. For comparison, the change in

[30] To check whether the downward bias of the price elasticity using specification (2) is caused by endogeneity of the knowledge stock, I re-ran the regression using specification (2) using a time trend as an instrument for the knowledge stock. Although a Hausman test fails to reject the null hypothesis that the knowledge stock is exogenous using specification (2) (the Wald statistic is just 0.6), the results from regressions using do provide evidence that increased patenting activity during the energy crisis biases the price elasticities downward in specification (2), as the price elasticity increases by a factor of 5 when using instrumental variables.

[31] This figure uses the absolute value of changes from year to year.

patenting activity because of the average change in prices is about 0.3 percent using specification (2), and 3.6 percent using specification (3).[32] Even during the peak of the energy crisis, the effect of energy prices ranges from a 0.5 to 5.5 percent increase in patents.[33]

Furthermore, it can be shown that omitting the stock of knowledge from the regressions biases the results. As Figure 1 shows, patenting activity in the energy fields increases quickly when energy prices rise but begins to fall before energy prices fall. Since the productivity estimates tend to fall over time, newer patents contribute less to the knowledge stocks than older ones. As a result, the values of the stocks themselves also fall over time. Since the quality of knowledge is important to future research, as suggested by the positive coefficients on the knowledge stock, diminishing returns to research help 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 hypothesis is confirmed by Table 5, which omits the knowledge stock from the regression. The first column presents regression results with no control for the existing knowledge stock. Without this control, both the elasticity of energy prices and government R&D become negative, although insignificant. Also, the mean and median lag is unrealistically high.

Next, I ask whether using patent data to construct stocks of knowledge is necessary or whether the returns to research can be adequately captured by a time trend. The second column of Table 5 provides the results of regressions in which the knowledge stocks from the preceding

---

[32] This calculation is simply the average change in energy prices during the sample period (5.78 percent) multiplied by the long run price elasticity.

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

section are replaced by a time trend. The results show the value of including the patent data, as once again the coefficients for energy prices and government R&D become negative.

Including the knowledge stocks is important because the rate at which the quality of knowledge falls over time varies across groups. For some groups, the quality of knowledge levels off during the 1980's, whereas for others the quality continues to fall. In fact, for fuel cells, the knowledge stock rebounds during the 1980's. Fuel cells experienced heavy patenting activity 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 existing stock of knowledge concerning new innovation, coupled with evidence of diminishing returns to research within the energy field, has important implications for dynamic models of environmental policy. The price elasticities that were found suggest that the reaction of the research community to a change in policy, such as a carbon tax, will be swift and that higher prices would quickly lead to a shift toward environmentally friendly innovation. However, since there are diminishing returns to research in a given field, firms will shift their research toward more productive areas as the marginal productivity in specific research fields declines. Thus, the burst of patenting activity resulting from a policy change is likely to be short-lived.

The results also show that diminishing returns are a factor in the falling patent-to-R&D ratio. Recall that one of the benefits of using a percentage of successful patent applications as the dependent variable in the induced innovation regression — equation (1) — is that it controls for the effect of exogenous factors on the propensity to patent, such as changing patent laws. Even with such exogenous factors controlled, there is still variation in the data that is explained by downward trends in the productivity of knowledge.[34] Also, note that many of the productivity estimates show an upswing near the end of the 1980's. This finding supports the recent results found by Kortum and Lerner (1998), who argue that increasingly fertile technology caused a surge in patenting in the United States during the 1990's.

B. Federal Energy R&D Spending

Interpreting the effect of government R&D spending on patents is difficult. Although the estimated elasticity of patenting to government sponsored R&D is very low, the results could be affected by changes in the emphasis of federally funded energy R&D, which changed after President Reagan took office in 1981. Before that time, 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 a program aimed at creating synthetic fuels from coal. When supporting research aimed at marketable technologies, federally funded energy R&D could be a substitute for private innovation. After Reagan's election, government funding for energy R&D was cut significantly. Department of Energy (DOE) support for research was limited to long-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 was expected that private firms would continue the R&D process by developing commercially acceptable products (U.S. Department of Energy 1987). As such, federally funded R&D performed after Reagan took office should be a complement to private innovation.

---

[34] This result is examined more thoroughly in Popp (1998).

To test the effect of changes in government R&D policy, I re-estimated equation (1), using specification (3) of the knowledge stock (productivity estimates only), with an additional term for government R&D done since 1981. The new regression is:

$$\text{(1')} \qquad \log\left(\frac{EPAT_{i,t}}{TOTPAT_t}\right) = \varphi_i + \gamma(1-\lambda)\log P^*_{E,t} + \theta \log K_{i,t-1} + \boldsymbol{\eta}(1-\lambda)\log \mathbf{Z}^*_{\mathbf{i,t}} + \chi(1-\lambda) d * R^*_{i,t} + \lambda^t \mu^0 + \varepsilon_{it}$$

where $R_{i,t}$ represents government R&D spending related to technology i in year t and d is a dummy variable equal to 1 for years greater than or equal to 1981. Table 6 presents the results of this regression. Note that the effect of government R&D performed since 1981 on patenting activity is over four times greater than government R&D performed before 1981, suggesting that government research after Reagan took office did become more basic in nature.

Despite the increased effect of government R&D after 1981, the magnitude of the overall effect of government R&D is small in both periods. Before 1981, federal R&D spending would need to increase by nearly $80 million to induce a new patent. After 1981, government R&D would need to increase by $15 million to induce a new patent. To put this result in perspective, consider that the average R&D expenditure per new patent is $1 million. Thus, the level of federal R&D spending needed to induce new energy patents is high relative to the level of spending done at the private level.

IV. Conclusions

I have used patent data to study the impact of energy prices on energy-saving technology. I have attempted to add 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 is the strong, positive impact energy prices have on new innovations. This finding suggests that environmental taxes and regulations not only reduce pollution by shifting behavior away from polluting activities but also encourage the development of new technologies that make pollution control less costly in the long run. My results also make 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 new innovation.

My results also shed light on the methodology for analyzing the determinants of technological change. In particular, I have shown 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 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 have demonstrated that patent citations can be used as a measure of the supply of knowledge available to inventors when they engage in research. Finally, the productivity estimates obtained from the citation data exhibit a downward trend, which suggests diminishing returns to energy research over time. Future research should address what factors influence the productivity of R&D, and whether policy can help to mitigate this downward trend.

REFERENCES

Ahmad, Syed, "On the Theory of Induced Innovation," Economic Journal, June 1966, 76(302), pp. 344-57.

Binswanger, Hans P. "A Microeconomic Approach to Innovation," Economic Journal, December 1974, 84(336), pp. 940-58.

_____. "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, 1978a, pp. 91-127

_____. "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, 1978b, pp, 128-63.

Caballero, Ricardo J. and Adam B. Jaffe. "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, 1993, pp. 15-74.

Cohen, Linda R. and Noll, Roger G. The Technology Pork Barrel, Washington, DC: The Brookings Institution, 1991.

Evenson, Robert E. "Patent Data by Industry," Cowles Discussion Paper no. 620, Yale University, 1991.

Goulder, Lawrence H. and Schneider, Stephen H. "Induced Technological Change and the Attractivenes of CO_2 Emissions Abatement," Resource and Energy Economics, August 1999, 21(3-4), pp. 211-53.

Greene, William H. Econometric Analysis, New York: Macmillan Publishing Co. 1993.

Griliches, Zvi. "Patent statistics as economic indicators: a survey," Journal of Economic Literature, December 1990, 28(4), pp. 1661-707.

_____. "Patents: Recent trends and puzzles," Brookings Papers on Economic Activity: Microeconomics, 1989, pp. 291-330.

Hicks, John R. The Theory of Wages, 1932, London: Macmillan.

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., Henderson, Rebecca and Trajtenberg, Manuel. "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations," Quarterly Journal of Economics, August 1993, 108(3), pp. 577-98.

Jaffe, Adam B. and Palmer, Karen. "Environmental Regulation and Innovation: A Panel Data Study," Review of Economics and Statistics, November 1997, 79(4), pp. 610-19.

Jaffe, Adam B. and Trajtenberg, Manuel. "Flows of Knowledge From Universities and Federal Labs: Modeling the Flow of Patent Citations Over Time and Across Institutional and Geographic Boundaries," 1996, Proceedings of the National Academy of Sciences, 93, pp. 12671-77.

Kamien, Morton I. and Schwartz, Nancy L. "Optimal Induced Technical Change," Econometrica, January 1968, 36(1), pp. 1-17.

Kortum, Samuel and Lerner, Josh. "Stronger Protection or Technological Revolution: What is Behind the Recent Surge in Patenting?" Carnegie-Rochester-Conference-Series-on-Public-Policy, June 1998, (48), pp. 247-304.

Lanjouw, Jean O. and Mody, Ashoka. "Innovation and the International Diffusion of Environmentally Responsive Technology," Research Policy, 1996, 25, pp. 549-71.

Lanjouw, Jean O. and Shankerman, Mark. "The Quality of Ideas: Measuring Innovation with Multiple Indicators," NBER Working Paper #7345, 1999.

Levin, Richard C., Klevorick, Alvin K., Nelson, Richard R. and Winter Sidney G. "Appropriating the Returns from Industrial Research and Development," Brookings Papers on Economic Activity, 1987, (3), pp. 783-820.

Mowrey, David C. and Rosenberg, Nathan. "The Influence of Market Demand Upon Innovation: A Critical Review of Some Recent Empirical Studies," Research Policy, 1979, 8, pp. 103-53.

Newell, Richard G., Jaffe, Adam B. and Stavins, Robert N., "The Induced Innovation Hypothesis and Energy-Saving Technological Change," Quarterly Journal of Economics, August 1999, 114(3), pp. 941-75.

Nordhaus, William. "Modeling Induced Innovation in Climate-Change Policy," Mimeo, Yale University, 1999.

Nordhaus, William. "Some Skeptical Thoughts on the Theory of Induced Innovation," Quarterly Journal of Economics, May 1973, 87(2), pp. 208-19.

Office of Technology Assessment. Energy Technology Choices: Shaping Our Future, OTA-E-493, Washington, DC: U.S. Government Printing Office, July 1991.

Popp, David. "The Productivity of Energy Research," presented at the 1998 NBER Summer Institute on Productivity, Working Paper #1998-4, The University of Kansas, 1998.

Public Citizen. "Renewable Energy Research and Development: An Alternative Budget Proposal for FY 1993-1995," Washington, DC: Public Citizen, 1992.

Rosenberg, Nathan. Exploring the Black Box: Technology, Economics, and History, Cambridge, UK: Cambridge University Press, 1984.

_____. Inside the Black Box: Technology and Economics, Cambridge, UK: Cambridge University Press, 1982.

Scherer, F.M. "Firm Size, Market Structure, Opportunity, and the Output of Patented Inventions," American Economic Review, December 1965, 55(5), pp. 1097-125.

Schmookler, Jacob. Invention and Economic Growth, Cambridge, MA: Harvard University Press, 1966.

U.S. Department of Energy, Energy Conservation: Multi-Year Plan FY 1989-93, Washington, DC: U.S. Government Printing Office, 1987.

Figure 1 – Patent Applications by U.S. Inventors

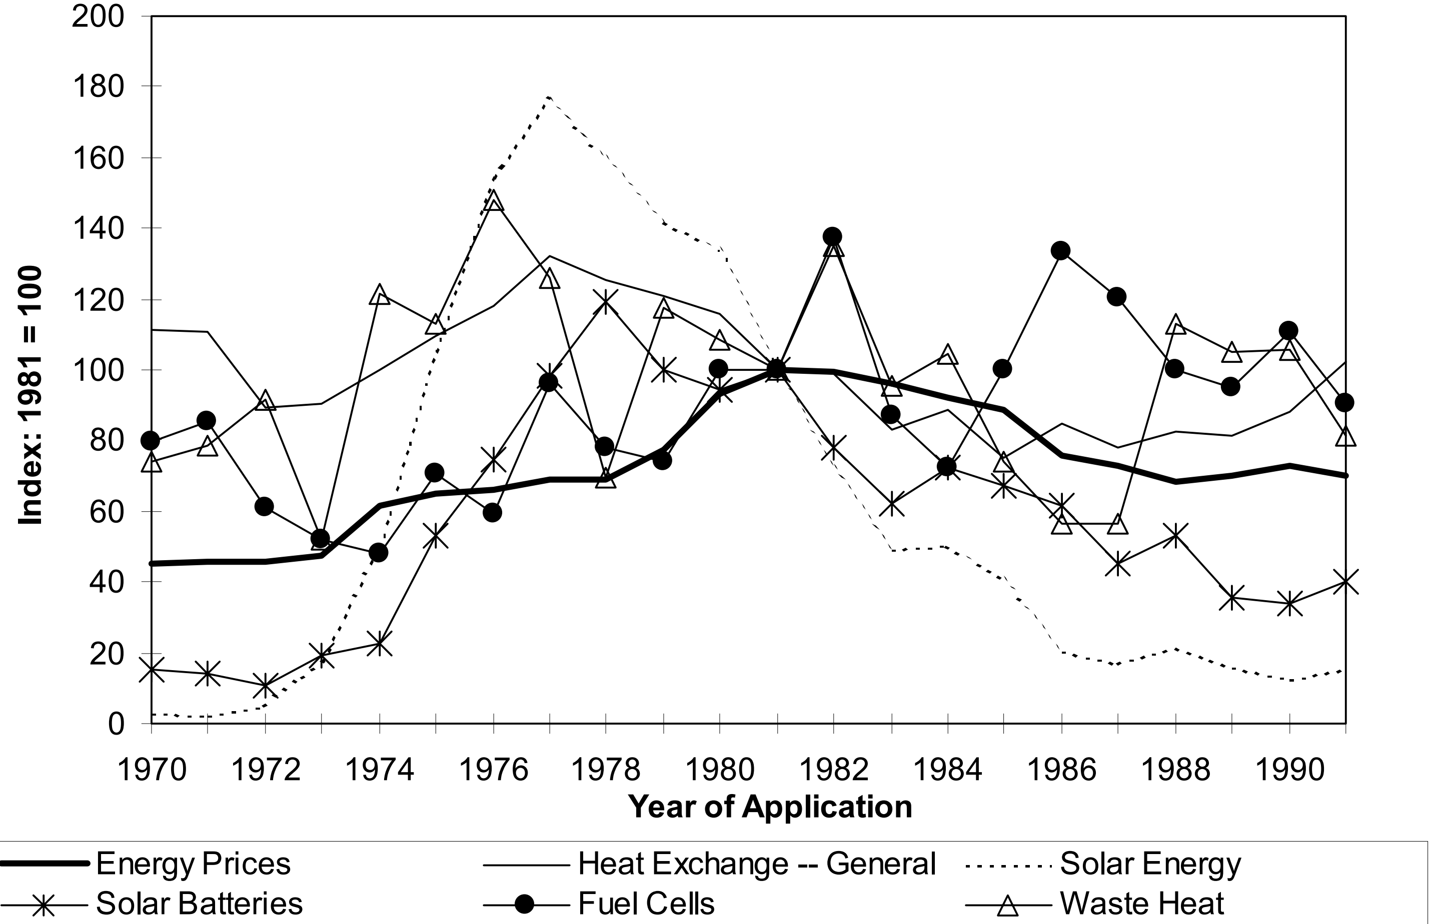

NOTES: Energy prices are the cost per million Btu of energy consumption, in 1987 dollars

Figure 2 – Probability of Citation

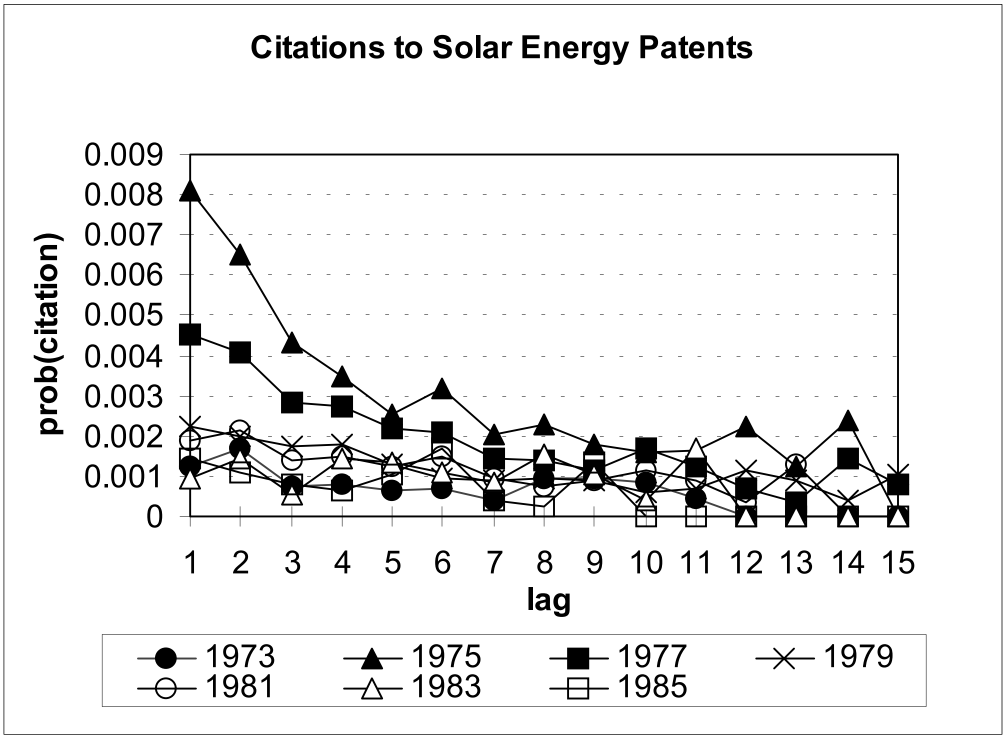

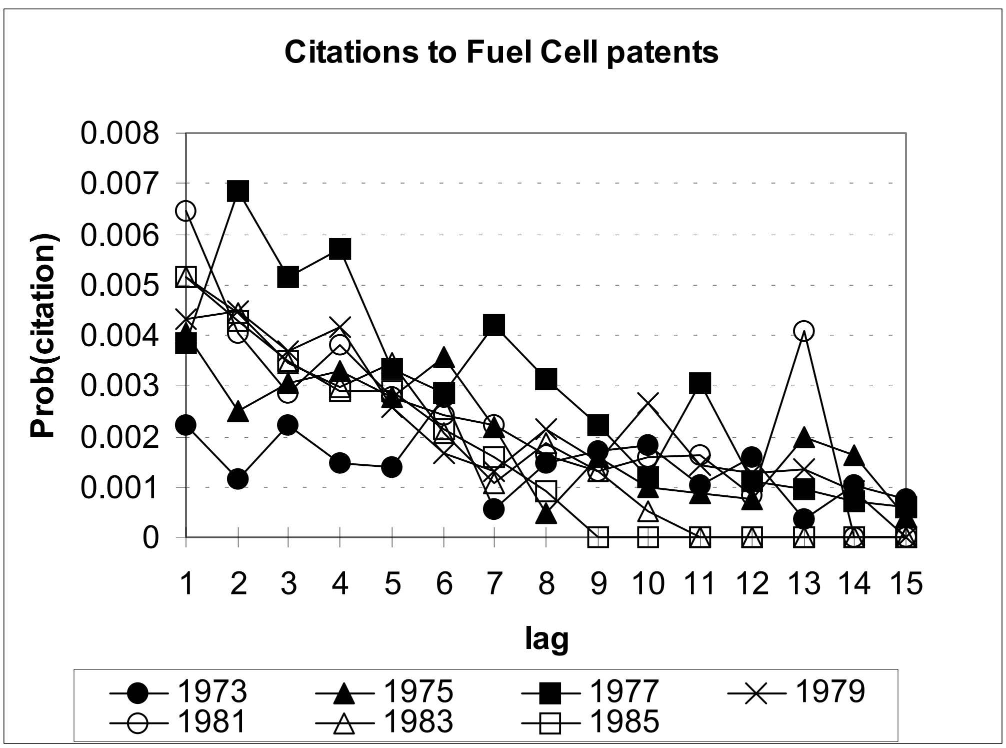

This figure presents the probability that patents granted in year *x* will be cited by patents applied for in following years. Each line represents the patents granted in a different year. The *x*-axis is the number of years since the patent was granted.

Figure 3 – Productivity Estimates

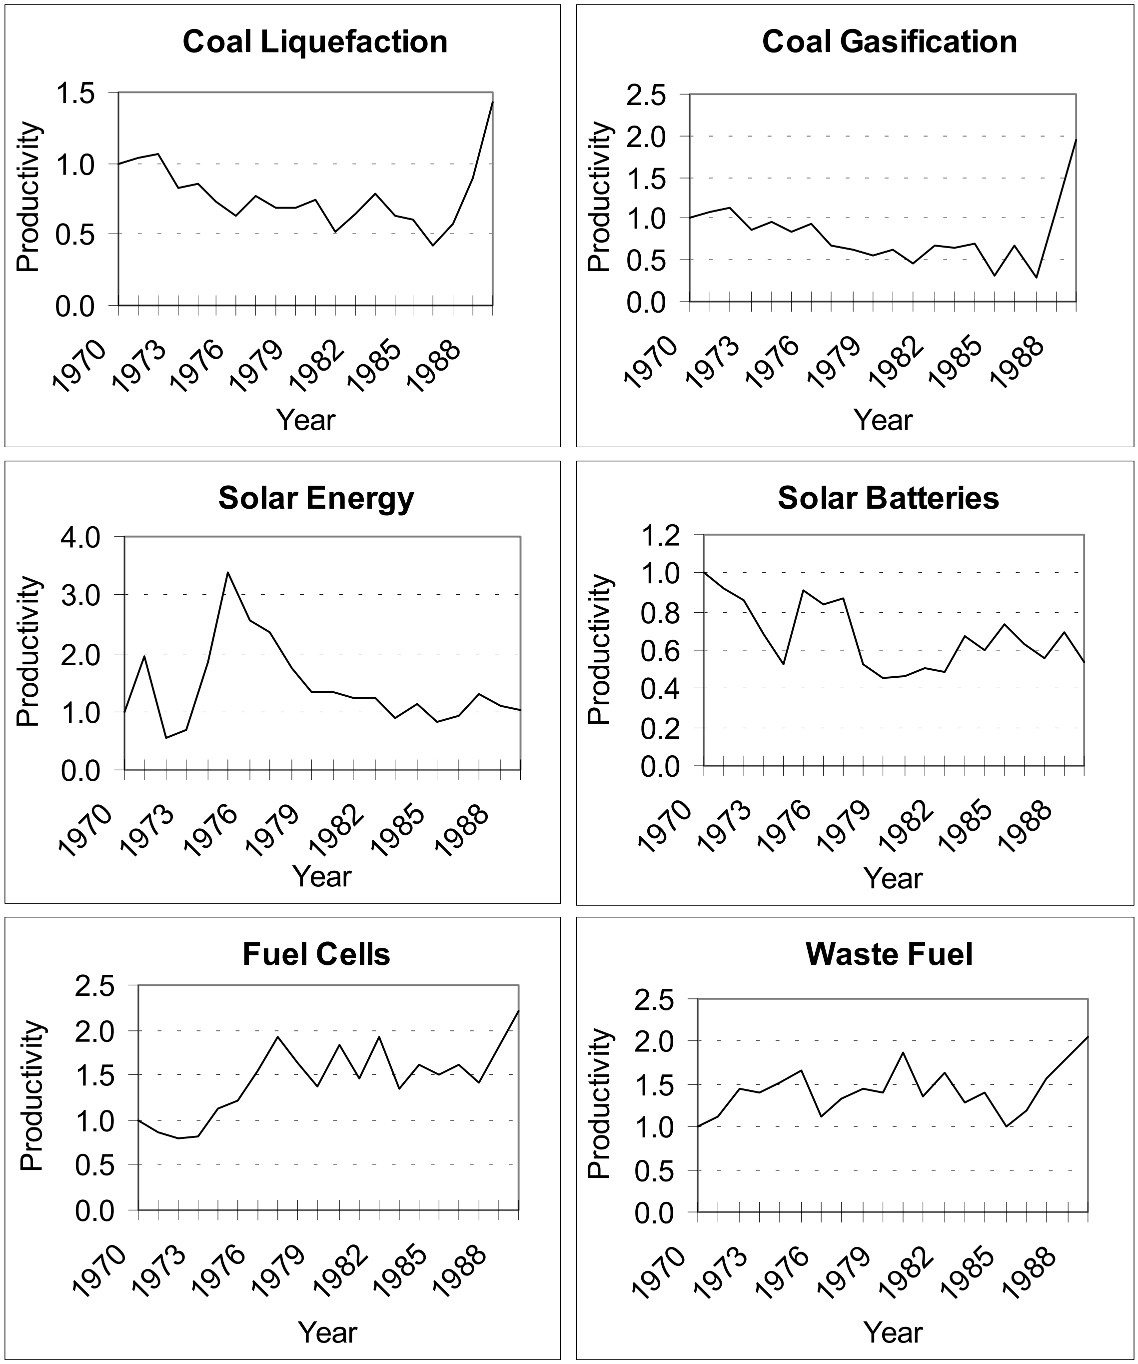

Figures show the productivity estimates for each technology group, with 1970 normalized to 1 in each case. Note that, for most technologies, there is a declining trend to the estimates, suggesting diminishing returns to research over time.

Figure is continued on the next page

Figure 3 – Productivity Estimates (continued)

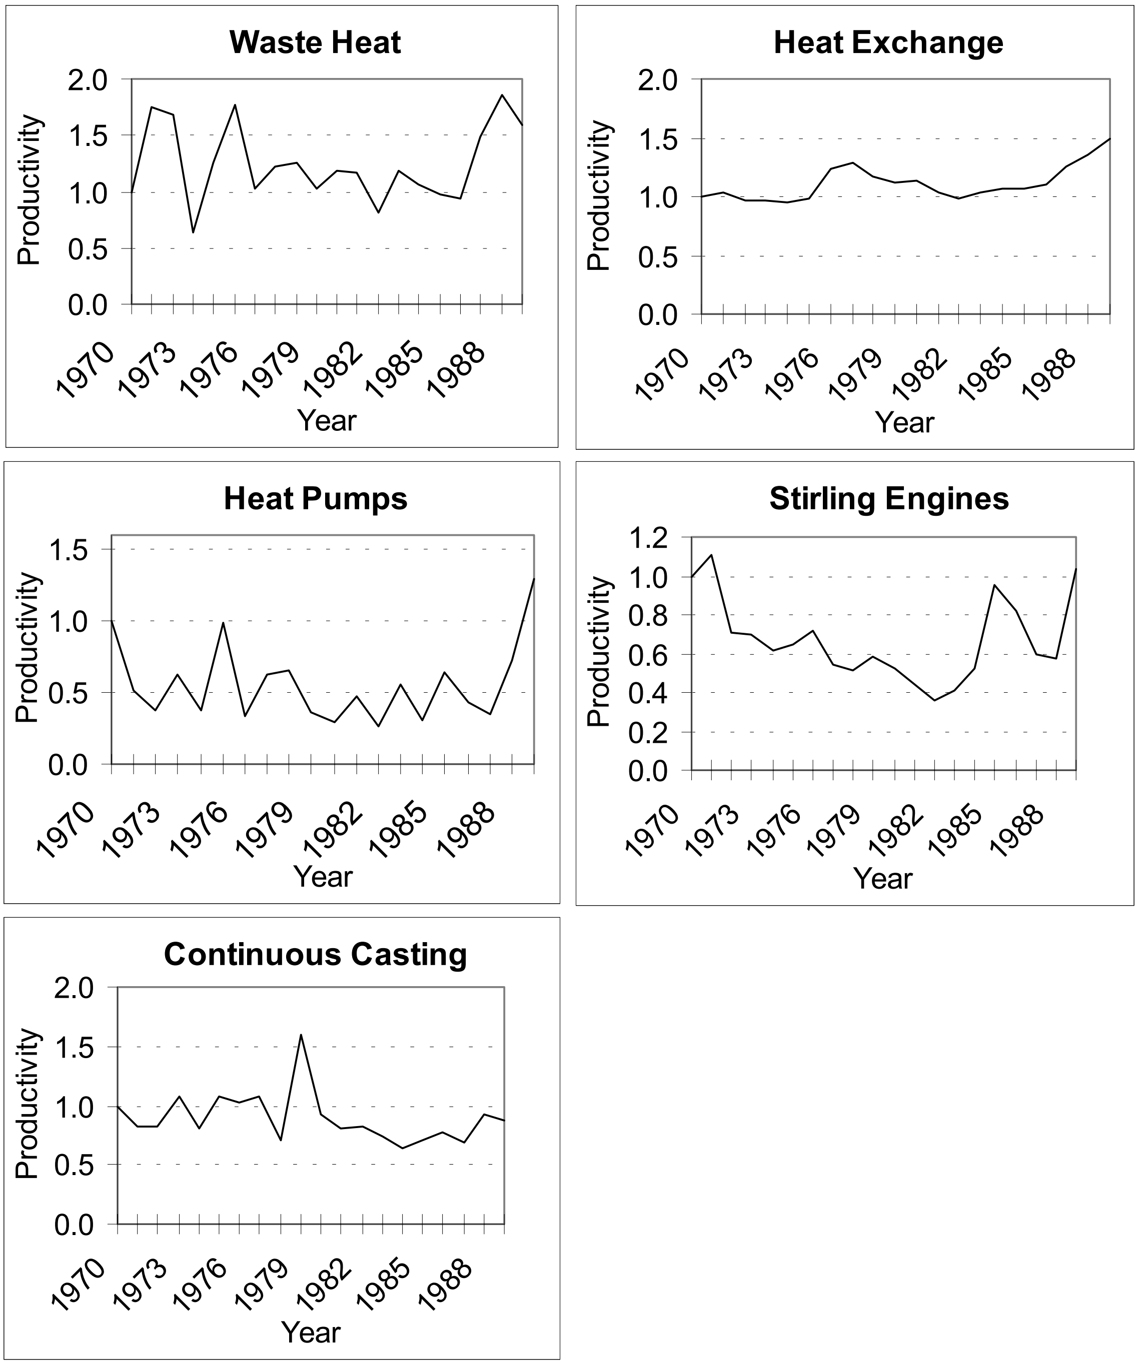

Figures show the productivity estimates for each technology group, with 1970 normalized to 1 in each case. Note that, for most technologies, there is a declining trend to the estimates, suggesting diminishing returns to research over time.

Figure 4 – Knowledge Stocks

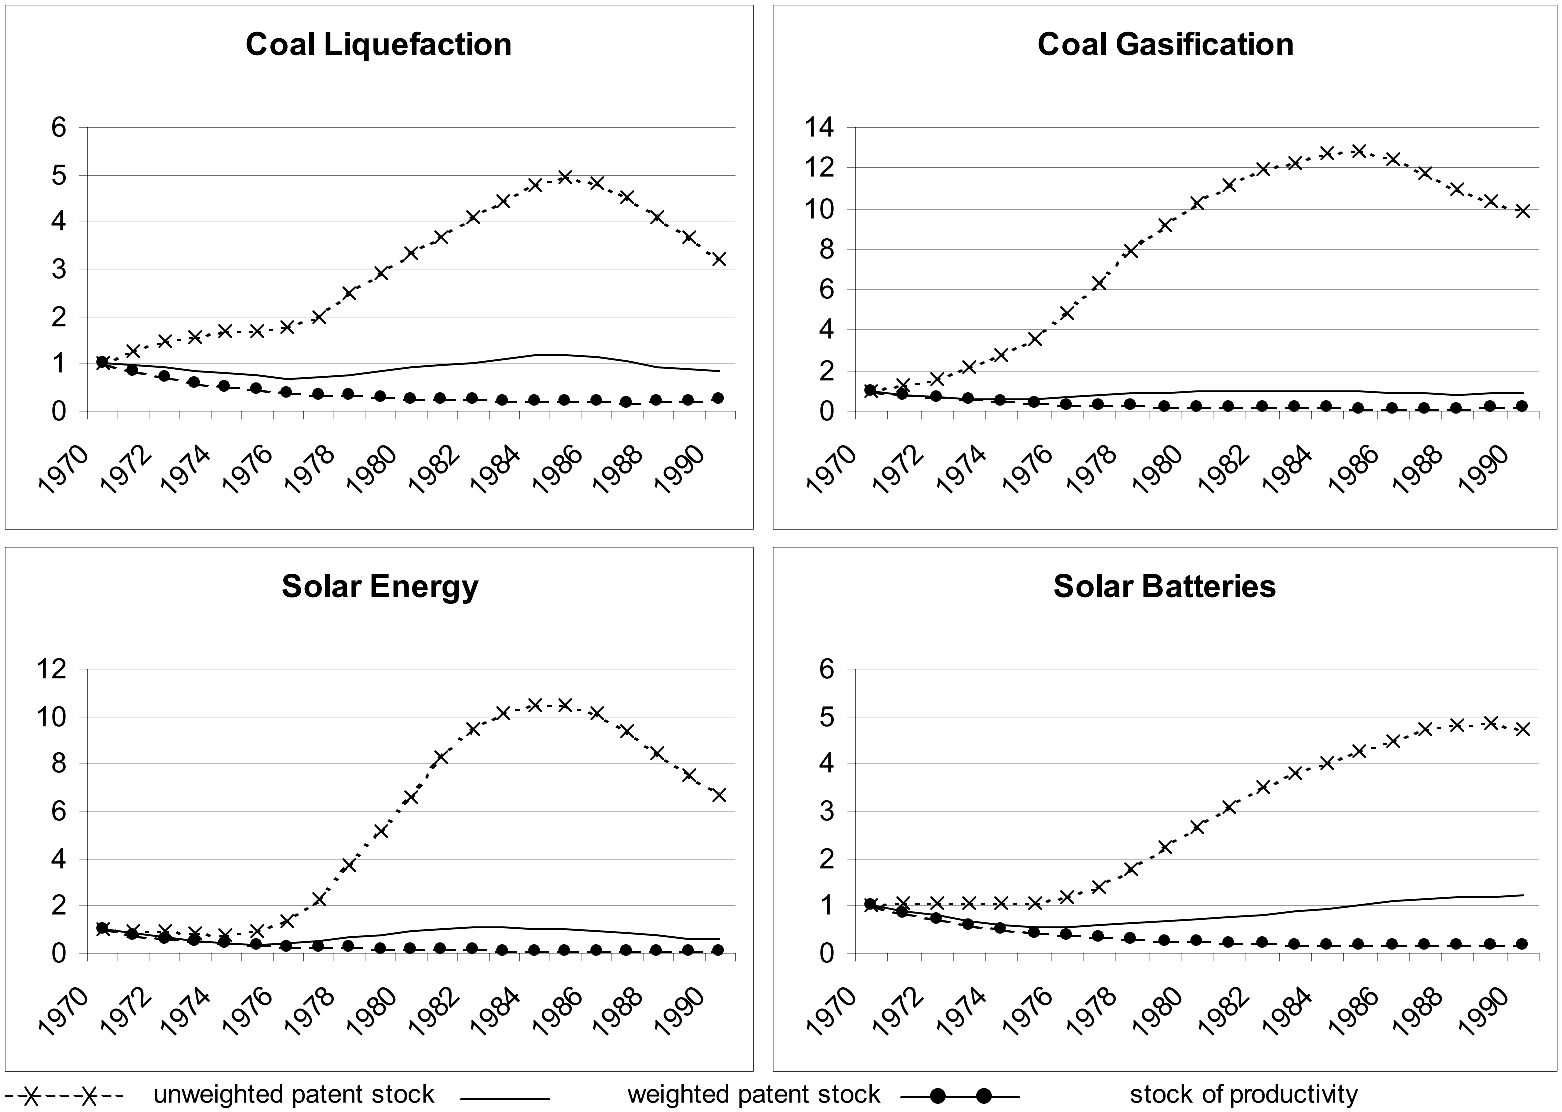

Figure 4 – Knowledge Stocks (continued)

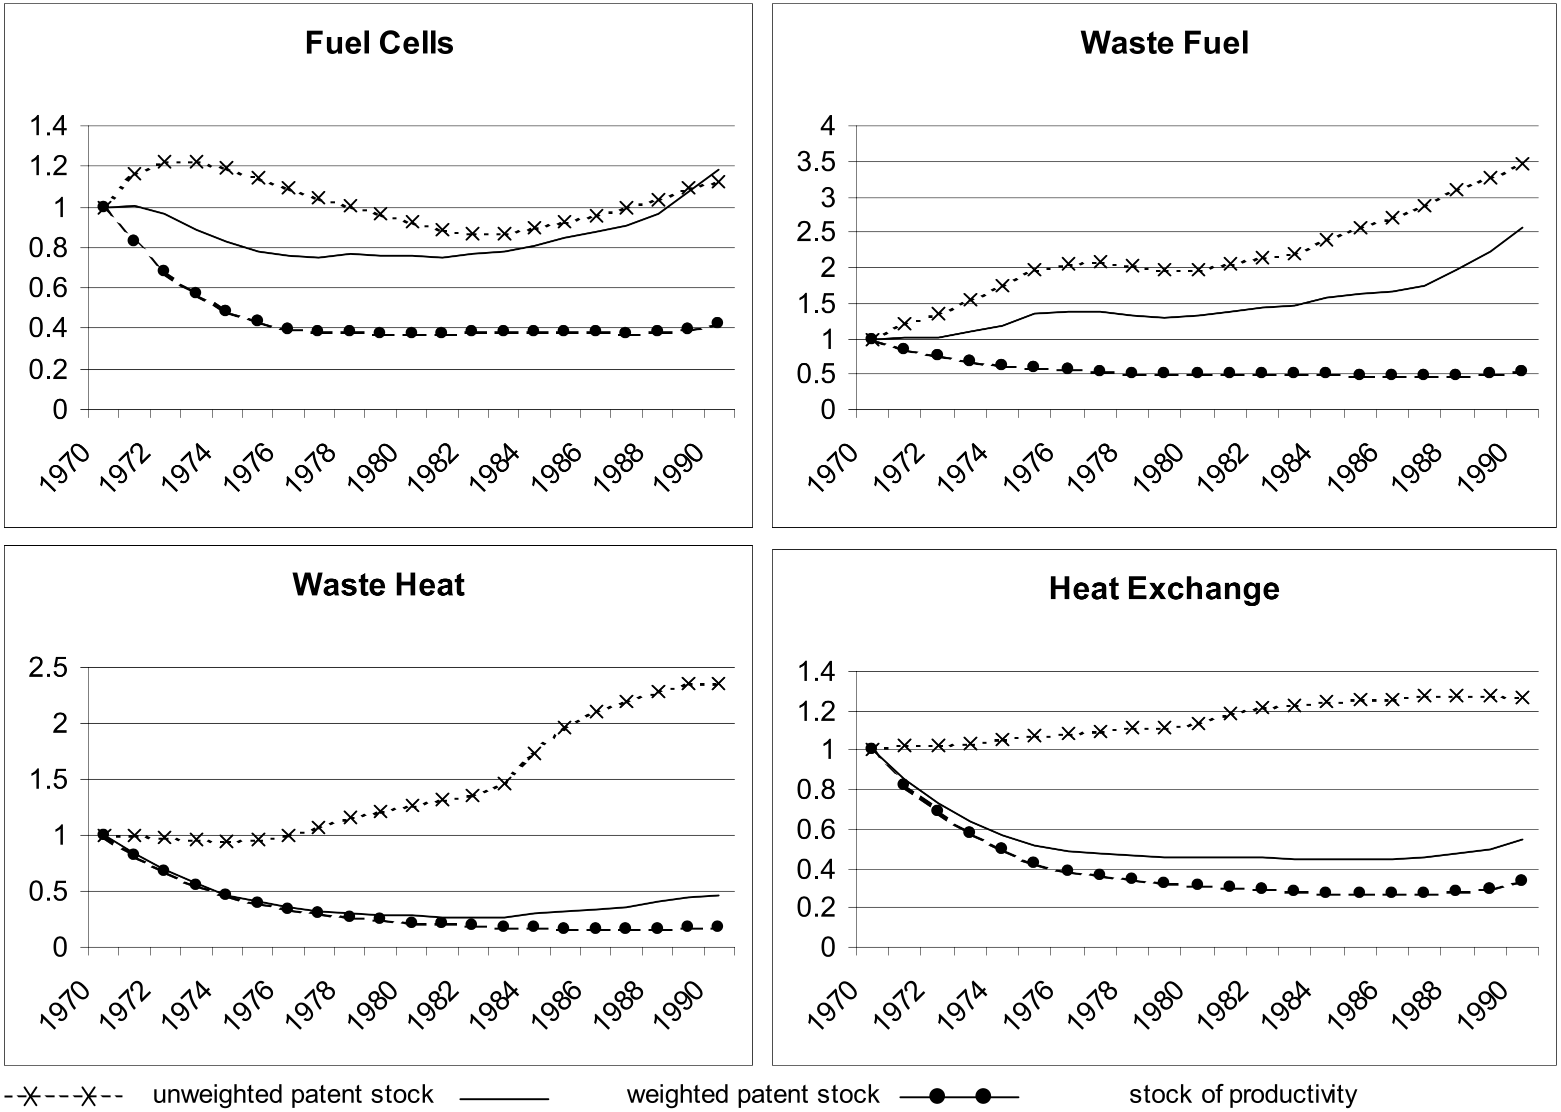

Figure 4 – Knowledge Stocks (continued)

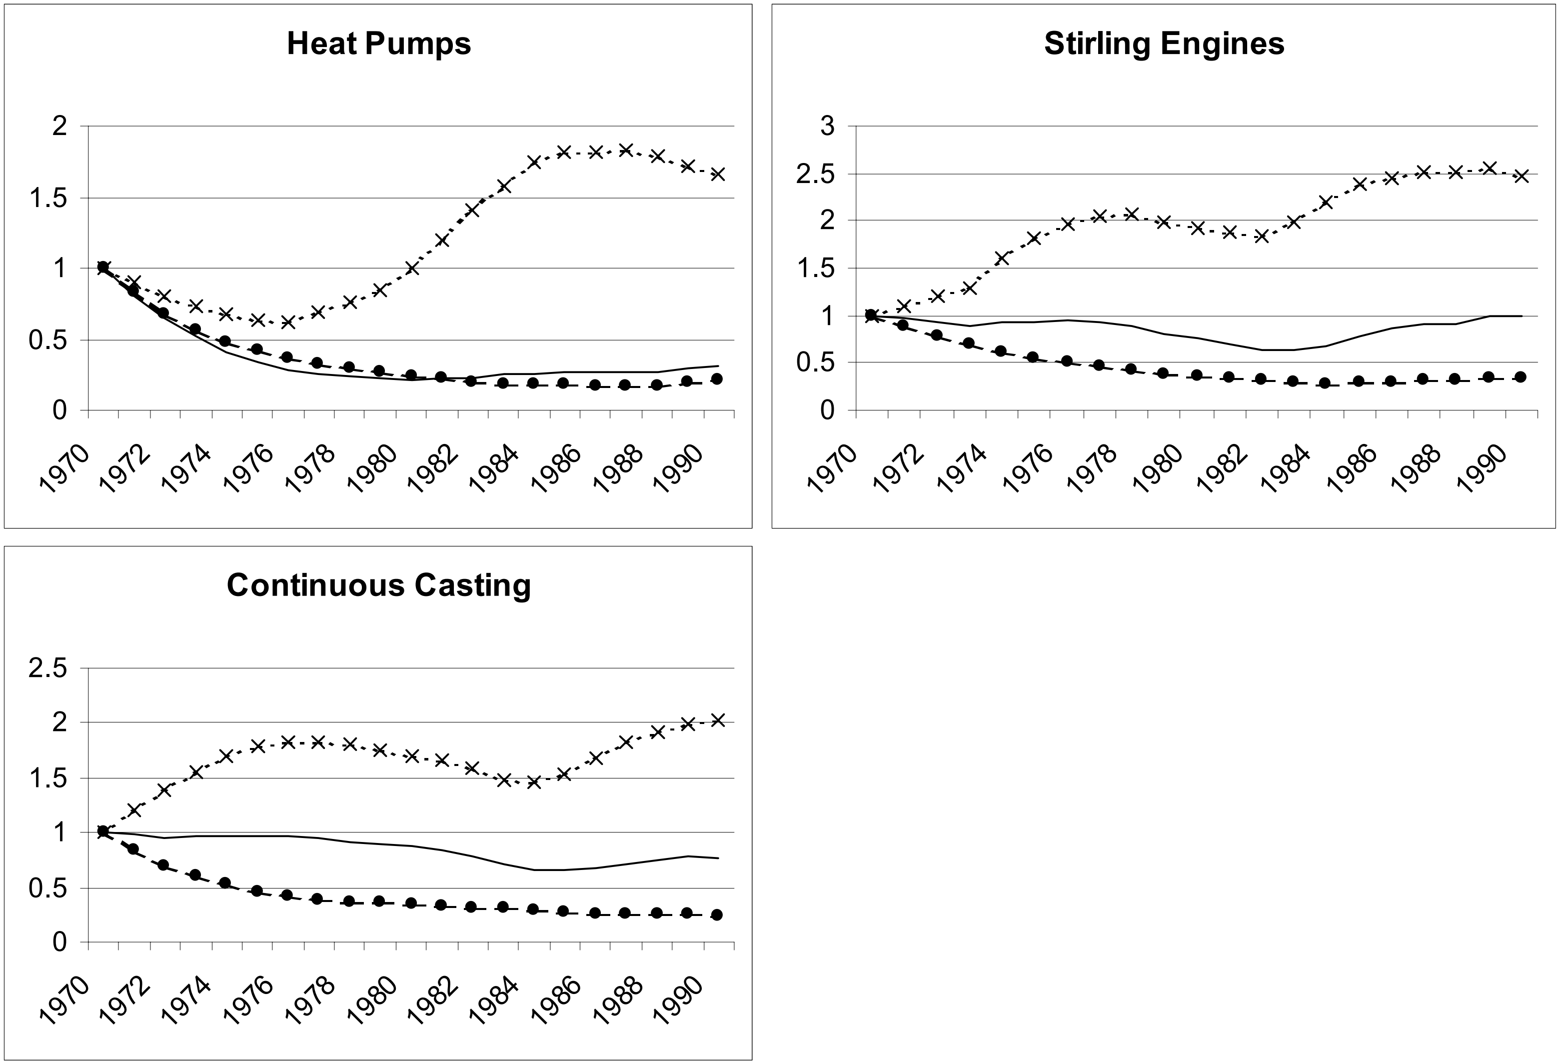

unweighted patent stock   weighted patent stock   stock of productivity

Table 1A – Summary Patent Data – Supply Technologies

Privately held U.S. patents sorted by year of application

Technology Group	**1970**	**1971**	**1972**	**1973**	**1974**	**1975**	**1976**	**1977**	**1978**	**1979**	**1980**	**1981**
Supply Technologies												
Coal liquefaction: producing liquid fuels	42	37	27	28	51	45	107	88	114	77	97	100
Coal gasification: producing gaseous fuels	14	24	16	20	38	31	42	45	53	32	38	27
Solar energy	6	5	10	36	104	218	321	367	333	295	278	208
Batteries for storing solar energy	18	17	13	23	27	63	89	117	142	119	112	119
Fuel cells	43	46	33	28	26	38	32	52	42	40	54	54
Using waste as fuel	63	53	52	52	49	29	32	34	41	40	50	44
Demand Technologies												
Recovery of waste heat for energy	17	18	21	12	28	26	34	29	16	27	25	23
Heat Exchange: general	425	423	340	346	382	418	450	505	479	462	443	382
Heat pumps	0	2	8	4	7	8	20	17	32	24	21	30
Stirling engines	13	13	12	9	17	11	17	11	12	11	18	21
Continuous casting processing of metal	84	115	67	63	48	46	43	37	40	45	44	43

Table 1B – Summary Patent Data – Demand Technologies

Privately held U.S. patents sorted by year of application

Technology Group	**1982**	**1983**	**1984**	**1985**	**1986**	**1987**	**1988**	**1989**	**1990**	**1991**	**1992**	**1993**
Supply Technologies												
Coal liquefaction: producing liquid fuels	82	74	70	34	20	12	14	22	16	10	12	8
Coal gasification: producing gaseous fuels	25	22	15	18	10	16	10	14	9	4	5	2
Solar energy	151	102	104	85	42	35	44	33	26	32	27	23
Batteries for storing solar energy	93	74	86	80	73	54	63	42	41	48	53	20
Fuel cells	74	47	39	54	72	65	54	51	60	49	61	58
Using waste as fuel	58	50	44	46	61	83	69	84	102	98	93	60
Demand Technologies												
Recovery of waste heat for energy	31	22	24	17	13	13	26	24	24	19	25	22
Heat Exchange: general	377	317	338	286	323	297	315	311	337	391	428	350
Heat pumps	18	11	8	14	15	11	5	14	18	22	14	17
Stirling engines	30	21	19	13	13	19	10	12	18	11	12	5
Continuous casting processing of metal	49	61	62	46	80	39	58	38	33	38	33	31

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

	correlation with:		
Technology	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
Waste as Fuel	-0.073	0.028	0.162
Waste Heat	0.283	0.055	-0.151
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
Weighted Average	0.077	-0.021	-0.122

Table presents correlations between energy prices and the number of patent applications per year.

Only privately held patents by American inventors are included. The average is weighted by the average number of patents in each technology group.

Table 3 – Miscellaneous Citation Regression Results

Probability of a patent being cited by a patent from:

citing years:	estimate:	standard error:
1974-1975	1.000	N/A
1976-1977	1.052	0.039
1978-1979	1.100	0.064
1980-1981	1.055	0.085
1982-1983	1.066	0.111
1984-1985	1.091	0.140
1986-1987	1.046	0.160
1988-1989	1.010	0.179
1990-1991	0.914	0.184

Group dummies:

technology group:	estimate:	standard error:
coal liquefaction:	4.364	0.750
coal gasification:	3.558	1.106
solar energy:	0.708	0.256
solar batteries:	2.272	0.466
fuel cells:	1.347	0.289
waste fuel:	1.511	0.417
waste heat:	1.427	0.793
heat exchange:	0.170	0.044
heat pumps:	9.212	2.157
Stirling engines:	6.222	1.266
continuous casting:	1.000	N/A

Rates of decay and diffusion of knowledge

	estimate:	standard error:
decay (β_1):	0.353	0.013
diffusion (β_2):	0.00199	0.00034

adjusted R-square: 0.755

Table 4 – Induced Innovation Regression Results

Dependent variable: % total domestic patent applications in each technology group

Independent Variables	unweighted stock of patents	weighted stock of patents	stock of productivity estimates
Constant	-9.015	-6.962	-6.631
	(-12.362)	(-29.097)	(-17.145)
Energy Prices	0.028	0.012	0.125
	(2.146)	(0.900)	(9.553)
Lagged Knowledge Stock*	0.719	0.834	0.439
	(25.612)	(66.695)	(6.521)
Government R&D	0.006	-0.002	0.003
	(0.968)	(-0.482)	(1.820)
Truncation error	1.924	-1.563	-0.026
	(2.445)	(-6.084)	(-0.092)
Lambda	0.933	0.781	0.800
	(18.905)	(10.191)	(21.389)
Long run energy elasticity	0.421	0.056	0.625
Long run govt. R&D elas.	0.085	-0.009	0.015
Median lag	13.81	3.57	4.00
Mean lag	9.92	2.81	3.11
GMM criterion	86.560	90.457	85.465
number of technology groups:	11	11	11

** -- Time trend and lagged values of other exogenous variables used as instrument for unweighted knowledge stock*

t-statistics below estimates
data from 1971-1991

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: % total domestic patent applications in each technology group

Independent Variables	no control for productivity	time trend
Constant	76.961	-101.147
	(0.646)	(-1.901)
Energy Prices	-0.116	-0.241
	(-0.708)	(-3.211)
Time Trend		23.477
		(2.005)
Government R&D	-0.001	-0.014
	(-0.554)	(-1.977)
Truncation error	-85.033	-7.394
	(-0.714)	(-1.936)
Lambda	0.996	0.958
	(181.489)	(52.226)
Long run energy elasticity	-30.429	-5.725
Long run govt. R&D elas.	-0.234	-0.331
Median lag	262.16	22.75
Mean lag	182.06	16.12
GMM criterion	61.226	85.484
number of technology groups:	11	11

t-statistics below estimates
data from 1971-1991

Table shows the results of induced innovation equations without controlling for the knowledge stock. Lagged party of the President and lagged government R&D used as instruments for government R&D.

Table 6 – Changes in Emphasis of Government R&D

Dependent variable: % total domestic patent applications in each technology group

Independent Variables	stock of productivity estimates
Constant	0.338
	(0.873)
Energy Prices	0.249
	(5.063)
Lagged Knowledge Stock	2.185
	(19.970)
Government R&D	0.033
	(4.845)
Government R&D*1981 dummy	0.136
	(6.394)
Truncation error	-6.444
	(-0.803)
Lambda	0.140
	(0.746)
Long run energy elasticity	0.289
Long run govt. R&D elas.	0.039
Post-1981 LR govt. R&D elas.	0.197
Median lag	0.16
Mean lag	0.35
GMM criterion	94.256
number of technology groups:	11

t-statistics below estimates
data from 1971-1991

Table shows the results of induced innovation equations allowing for differences in the effect of government R&D before and after 1981. The stock of knowledge is a stock of productivity estimates only. Lagged party of the President and lagged government R&D used as instruments for government R&D.