PRP
7/28

The Productivity of Energy Research

David Popp*

May 19, 1998

How do the returns to research and development (R&D) evolve over time? I address this question by studying the productivity of energy research during the 1970's and 1980's. Patent citation data are used to obtain productivity parameters for 12 different energy technology groups. Evidence of diminishing returns to research is found: the productivity parameters fall over time, and are lower in years with more patents. Furthermore, the productivity estimates are found to coincide with a falling patent-to-R&D ratio in energy fields. Implications for the more general fall in the patent-to-R&D ratio in the U.S. economy as a whole are discussed.

* – Department of Economics, University of Kansas, 213 Summerfield Hall, Lawrence, KS, 66045-2113, dpopp@ukans.edu

The author would like to thank William Nordhaus, Joel Waldfogel, Robert Evenson, and Ariel Pakes for helpful comments. Despite their helpful suggestions, however, all mistakes are solely the responsibility of the author. In addition, the assistance of Robert Evenson, Jon Putnam, and Samuel Kortum in locating patent data is gratefully acknowledged. Finally, thanks to Adam Jaffe for providing the basic programming code for the estimation in section III. Partial financial support for this paper comes from DOE grant 593A3140217.

How do the returns to research and development (R&D) evolve over time? This fundamental question lies at the heart of several important questions in economics. Several studies have addressed the returns to R&D by looking at the value of the *output* of R&D. They do this either directly, by including R&D expenditures in the production function of a firm, or indirectly, by estimating the value of holding patent rights.[1] This paper approaches the question from a different direction. Rather than asking the value of the output of the research process, I ask what is the expected value of the *inputs* to the research process – that is, how much will additional R&D spending increase the expected value of the outcome of the R&D project. The expected value of R&D spending is an important determinant of the amount of R&D spending in any given time period. Empirical work on the determinants of R&D spending have been burdened by a lack of data on the research opportunities facing inventors. This paper uses patent citation data to fill this gap.

The pattern of returns to research over time is important for several reasons. This paper focuses on two questions for which the productivity of research inputs provide valuable information. The primary motivation of this research is that the productivity of research inputs is an important determinant of a firm's research investment decision. Many models of innovation make assumptions about the returns to R&D. For example, the induced innovation literature, which studies the impact of factor prices on innovative activity, postulates the presence of an innovation possibilities frontier (IPF), which represents the possible technological advances in a given period.[2] However, that literature does not successfully address how the opportunities for induced innovation will vary over time. Instead, the frontier is treated as exogenous. No relationship between past research and the current IPF is posited. Yet, unless returns to research are constant over time, some relationship must exist. This, in turn, will cause the inducement effect of demand variables to vary over time. As noted in Nordhaus (1973), the failure of the induced innovation literature to address the formation of the innovation possibilities frontier prevents the theory from truly endogenizing technological change. This paper addresses that criticism by examining how previous research efforts affect the productivity of future research.

A second reason for concern over the productivity of research inputs comes from data on patents and R&D spending. Over time, the ratio of patents to R&D expenditures has fallen in all industrial nations. Figure 1 illustrates this trend for the United States. In addition, total factor productivity growth declined during the 1970's. Noting the correlation between the declining patent-to-R&D ratio and total factor productivity growth, some researchers have attributed this fall to a decline in the productivity of research.

Hypotheses to explain the falling ratio focus either on the returns to R&D, such as Evenson (1991) or Kortum (1993), or on exogenous changes in the willingness to patent (Griliches 1989, 1990). Those who argue that exogenous changes in patenting behavior are the cause of the declining patent-to-R&D ratio argue that the returns to research have not fallen over time, but rather that inventors are less likely to patent successful innovations that they were earlier.[3] Changes in industry composition or legal rulings by the courts could change the perceived value of patent protection. By focusing only on the returns to R&D, this paper will shed some light on the puzzle of the falling patent-to-R&D ratio.

This study of the productivity of research R&D is carried out by examining trends in energy-saving innovations over the past 25 years. To show how the trends in research productivity vary across technologies, several different energy-efficient technologies are included. Studying energy-efficient innovations are of interest because of the insight they provide into the role that technological change can play in alleviating many of today's environmental problems. Many of these, such as global warming, are long-term problems. Technological change is likely to play a key role in alleviating them. In addition, several papers show that policymakers can induce environmentally-friendly innovation with policies such as a carbon tax or a regulation restricting emissions[4]. Knowing how the productivity of such research varies over time is important to understanding just how important a role technological change will play in easing environmental concerns. If evidence of diminishing returns is found, the effects of induced innovation from a policy change are likely to be short-lived.

This paper uses patent citation data as a means of quantifying the research opportunities available. When a patent is granted, the patent document includes citations to earlier patents that are similar to the new invention. The legal purpose of these citations is to limit the claim of the new patent – any part of the invention that duplicates things from the cited patents cannot be claimed by the holder of the new patent. Previous work with patent citations has focused on the flow of knowledge across institutions (such as from universities to private industry), across regions, and across nations (Jaffe and Trajtenberg, 1995; Jaffe, Henderson, and Trajtenberg, 1993). In addition, patent counts weighted by the number of subsequent citations have been found to be a good measure the social value of innovation (Trajtenberg, 1990). Finally, Caballero and Jaffe (1993) use citations in a macro growth model both to study the diffusion and obsolescence of knowledge, and to study the productivity of knowledge. The current paper is most similar in spirit to the Caballero and Jaffe work, in that the focus is on changing citation patterns over time. However, this paper takes a more micro-oriented approach by studying citations in several different technology areas.

The paper is organized as follows. It begins with a review of the existing literature on the productivity of R&D. Section II discusses how patent citations will be used as a measure of the productivity of research inputs. A model of research under uncertainty, first developed by Evenson and Kislev (1975), is introduced. Using the implications of their model, I discuss how patent citation data can be used as a proxy for the opportunities facing inventors. In section III, a model of patent citations, first developed by Adam Jaffe and his co-authors (Caballero and Jaffe, 1993; Jaffe and Trajtenberg, 1995), is used to estimate parameters for this opportunity. These are referred to as *productivity parameters*. The validity of these parameters is tested in section IV. Section V examines the behavior of the productivity estimates over time. Diminishing returns are found, both across time and with respect to the number of patents granted in a given year. The productivity estimates are compared to the falling patent-to-R&D ratio in energy research, and shown to explain the trend well. More general implications for the falling patent-to-R&D ratio in the economy as a whole are discussed. Section VI concludes.

I. Previous Studies of the Returns to R&D

This paper examines the productivity of research *inputs* – that is, what do inventors perceive as the expected value of research before a research project is begun. It uses subsequent citations to patents in the energy fields to make inferences about the potential for successful research at any given time. In contrast, several papers have examined the productivity of research *output* – that is, what are the contributions of the results of R&D to economic growth. Griliches (1995) presents a survey of this work.

Griliches outlines three approaches used to assess the effects of R&D on economic growth. The first, historical case studies, find high private and social rates of return on new inventions. However, only focusing on successful innovations may bias the results of such case studies. The second line of research is the use of invention counts and patent statistics. Much of this work focuses on patent renewal data. In most countries, an annual fee must be paid to keep a patent in force. Using data on the number of patents that are renewed and the number allowed to expire, researchers such as Pakes (1986) and Pakes and Schankerman (1984) have found that the value of patents are highly skewed. Most patents have little worth, but a few have enormous value. Lanjouw, Pakes, and Putnam (1996) provide a review of empirical work using patent renewal data. However, to interpret the results of renewal data studies, it is important to note that these studies capture the value of patent *protection* to the inventor, rather than the value to society of the invention itself.

The third line of research on R&D productivity estimates production functions that include R&D expenditures as a right-hand variable. Examples include Griliches and Jaques Mairesse (1984), Clark and Griliches (1984) and Scherer (1982, 1984). The rate of return to R&D found by these studies generally lies between 0.2 and 0.5, where the rate of return is defined as the change in the growth rate of output which results from additional R&D spending. Some evidence is found to suggest that research was less productive during the 1970's. However, Griliches notes that this change in productivity is not enough to explain the fall in total factor productivity growth during this period. In addition, the returns to R&D appear to recover by the end of the 1980's, bringing the hypothesis of diminishing returns to R&D into question. These studies are complicated by all of the econometric difficulties of estimating production functions, such as simultaneity, and also by data problems, as measures of output do not adequately account for changes in quality brought about by successful innovation.

In contrast with these papers, this study looks at the productivity of research *inputs*. To understand the distinction between the inputs and outputs of research, consider a production function of the form:

(1) $$Q_t = F_t[X_t(K_t(R_t, K_{t-1})].$$

Equation (1) states that output in time t, denoted Q_t, is a function of the inputs used, X_t. The productivity of the inputs is a function of the current state of knowledge, K_t. The current state of knowledge is a stock that depends on current research R_t, and the productivity of the previous stock of knowledge. The previous studies mentioned examine the productivity of the output of R&D – that is, they attempt to measure $\partial Q_t/\partial K_t = (\partial Q_t/\partial X_t)*(\partial X_t/\partial K_t)$. In contrast, this paper is concerned with the marginal benefits of additional research – that is, $\partial K_t/\partial R_t$. In particular, this paper looks at how past knowledge affects the productivity of current research, or $\partial^2 K_t/\partial R_t \partial K_{t-1}$.

Patent citations will be used for this purpose. The one previous study that has related patent citations to the returns to research is Caballero and Jaffe (1993). They use a database of U.S. patents from 1975 to 1992, and examine citations contained in these patents to patents dating back to 1900. Among their findings is that newer patents are less likely to be cited, suggesting a decline in the value of knowledge represented by these patents.

This paper builds on the work in Caballero and Jaffe. Rather than focusing on all inventions, we look only at inventions in the area of energy efficiency. By studying several different technologies within this area of research, I am able to see how research opportunities vary across fields. In addition, annual estimates of the productivity of past knowledge are obtained, so that changes in these estimates can be studied more carefully than in Caballero and Jaffe. The estimates are related to the amount of innovative activity in these fields. I find that the productivity estimates are significant predictors of the overall level of R&D in a given field. Finally, the pattern of productivity estimates across time is examined for diminishing returns.

II. Patent Citations and the Productivity of Earlier 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 and the patent examiner. It is the applicant's responsibility to list any related previous patents of which he or she is aware. In addition, the examiner, who specializes in just a few patent classifications, will add other patents to the citations, as well as subtracting any irrelevant patents cited by the inventor. Patent citations narrow the reach of the 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.[5] For the same reason, inventors have incentives to make sure that no unnecessary patents are cited. As a result, the previous patents cited on a new patent should be a good indicator of previous knowledge that was utilized by the inventor.

My interest in patent citation data is to develop a proxy for the expected value of research at any given time. Citations provide evidence of the existing state of technology when an invention is completed. Just as citations to earlier work in a journal article indicate that the cited works were of use to the author, citations to previous patents indicate that earlier innovations were of use to the inventor.[6] Frequent citations to a patent indicate that the knowledge embodied in that invention is useful to other inventors. When several past useful inventions are available to inventors, current research should be more productive.

To visualize this, it is useful to consider the R&D process of a firm more carefully. In this section, I summarize the search model of R&D (Evenson and Kislev, 1975).[7] The model provides an explicit formalization of the expected value of research. Although the data in this study are not sufficient to estimate the expected value of research as set out in the search model, the model does provide useful intuition as to how patent citation data can be used to infer this information.

The results of R&D are not certain when the investment is made. Acknowledging the uncertainty of the research process, Evenson and Kislev (1975) model scientific research as a process of searching over a distribution of possible outcomes. The distribution represents the value of the possible outcomes from a given level of research. The current level of knowledge is represented by one point in the distribution. Since the probability of successful research varies across different technologies, each technology is assumed to have its own distribution. For example, the invention of a new engine and the invention of a new solar collector presumably come from different distributions of potential outcomes.

Figure 2 presents a distribution of outcomes for a hypothetical research field. The x-axis measures the value of the invention that results from research. K_t represents the current level of knowledge. The search model portrays research as a series of random draws taken from this distribution of possible outcomes. A number of draws are taken each period, and the result from the highest draw is considered the output of the research that period. If the highest draw is greater than the previously existing knowledge, research is successful. As a result, the expected value of research is represented by the area to the right of K_t (shaded in Figure 2). This area can be thought of as the *stock of unexploited technology*. It tells us how useful the existing stock of knowledge (K_t) is to inventors. When K_t is further left in the distribution of possible outcomes, there are more opportunities for research that follows up on this idea. Thus, the larger the stock of unexploited technology, the greater the expected value of research.

The distribution of possible outcomes also provides insight as to how the returns to research can vary over time. If a successful invention occurs in period t, the state of knowledge in period t+1 shifts rightward. As a result, the expected value of research will be less in period t+1. Evenson (1991) refers to this as *invention potential exhaustion*. The intuition is that there is a limited pool of potential inventions on which to draw. As ideas in this pool are used up, the likelihood of future successes dwindles. Invention potential exhaustion is illustrated in Figure

3a. The expected value of research in time t is equal to A+B. An invention in period t shifts out the level of knowledge facing inventors in period $t+1$. As a result, the expected value of research in period $t+1$ is only area B.

However, there is a second possibility. The preceding discussion assumes that the distribution of potential outcomes does not change over time. This is not necessarily the case. Advances in science may change the nature of the distribution of possible outcomes. New innovations may contain basic knowledge that make possible in period $t+1$ inventions that were not possible in period t. Evenson refers to this as *recharge*. Recharge occurs when inventions shift the distribution of potential outcomes rightward. Recharge can also come from inventions in related fields or from advances in basic science. For example, the invention of microcomputer chips led to the development of a generation of electronic equipment.[8] Recharge is illustrated in Figure 3b. Again, the original state of knowledge is K_t, and the new state of knowledge is K_{t+1}. However, in this example, the distribution of potential outcomes has shifted outward. The expected value of research in period t is A+B; in period $t+1$ it is B+C. Since area C is drawn to be larger than area A, the effect of recharge dominates invention potential exhaustion. As a result, there are no diminishing returns to R&D in this example.[9]

As Evenson notes, the debate over returns to R&D centers on the possibilities for recharge. The principles of diminishing returns should apply to research just as much as any other economic activity, leading to invention potential exhaustion. *The question is whether new advances in science – or recharge – occur quickly enough to counteract the effect of exhaustion.* In this paper, I look at variations in the expected value of research over time to address this question.

As figures 2 and 3 show, two pieces of information would be needed to estimate the expected value of research as modeled by Evenson and Kislev:

1) the properties of the distribution of possible outcomes, and

2) the position of the current level of knowledge within this distribution.

Unfortunately, neither of these is observed directly in the data set used in this paper. However, Evenson's search model of R&D does provide insight as to how to make inferences about the expected value of research using patent citation data.

The expected value of current research is higher the farther left the current state of knowledge is in the distribution of possible outcomes. If the patent representing the current state of knowledge is at the left of a distribution, it should be followed by many improvements to it, and thus be cited often by subsequent patents. Conversely, if an invention is at the right of the distribution of potential outcomes, subsequent patents will cite it less often. By observing the subsequent citations to patents from any given year, one can infer how large the stock of unexploited technology was.

Figure 4 illustrates how the expected value of research relates to the nature of the distribution, and how citation data can be used to infer this information. Consider a firm considering two research projects, each represented by a different distribution of possible outcomes. K_{ti} represents the current state of knowledge in each field. Research is only successful if it results in a draw from the distribution greater than K_{ti}. The expected value of research is the area under the distribution and to the right of the current level of knowledge.[10] To simplify matters, assume that the only difference between the two distributions is the level of the current state of knowledge in each field. Figure 4 is drawn so that $K_{t2} < K_{t1}$. There is a greater stock of unexploited technology remaining to be discovered by the second project (represented by the shaded areas in figure 4).

Thus, by looking at the subsequent cites to the patents representing these inventions, one can infer where in the distribution of possible outcomes they lie, and thus how large the stock of unexploited technology is. To do this, I assume that inventions at the upper end of a distribution of potential outcomes will cite earlier inventions in the distribution. Of course, in any given year, there will be many inventions, taken from many different distributions. For example, the invention of a new engine and the invention of a new solar collector presumably come from different distributions of potential success. Inventions that are at the left of a distribution will be followed by many improvements to it, and will thus be cited often. Inventions at the right of a distribution will be cited less often. In figure 4, we see that there are more potential inventions available for the second project. Thus, if K_{t1} and K_{t2} represent previously patented innovations, we would expect more citations to the patent for K_{t2} than to the patent for K_{t1}.

III. Estimation

The previous section shows that a higher number of citations to a patent results in a higher expected value of R&D for future inventors. The goal of this section is to observe the likelihood that patents granted in any given year are cited by patents from subsequent years. However, 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, it is necessary to look at the probability of citation, rather than a pure count of citations. In addition, because most patents are never cited, patents are grouped by years, rather than using individual patents as data points.

To begin, cohorts of patents that could potentially cite each other are created. For each technology group, potentially cited patents are sorted by year of grant, and are denoted *CTD*. 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, so that the results of this paper will correspond to the results of Popp (1997, 1998), in which energy patents were sorted by year of application to estimate a model of induced innovation. Citing patents are denoted *CTG*. The data are sorted by *CTD,CTG* cohorts. Separate cohorts are constructed for each technology group, *i*. For example, one cohort might be citations to all solar energy patents granted in 1975 made by solar energy patents applied for in 1980. Denoting citations as *c* and the number of patents from each year as *n*, the probability of citation for patents within each cohort is:

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

The goal is to find the likelihood that patents granted in year *CTD* will be cited by patents in *any* subsequent year. To find this, it is necessary to control for other extenuating factors that affect the likelihood of citation. The model used in this paper builds on a model used in Caballero and Jaffe (1993) and Jaffe and Trajtenberg (1995) for that purpose. These authors identify three factors that affect the probability of a patent being cited by subsequent patents. They are:

1) the rate at which newly produced knowledge, as represented by a newly-patented innovation, diffuses through society,
2) the rate at which new innovations become obsolete, as they are replaced by better inventions, and
3) the ways in which citing behavior has changed over time.

In addition, a fourth factor, the size of the technology group, is added for this paper. 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 patent 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 patents in the group are less likely in these groups.

The reasons for considering the decay and diffusion of knowledge are straightforward. It is not the production of a new process in a laboratory that spurs further innovation, but rather the dissemination of the new knowledge throughout society. The rate of diffusion accounts for the delays in new knowledge becoming publicly available. This effect would cause older patents to be cited more than newer patents. Working against the diffusion process is the obsolescence of older knowledge. As time passes, newer inventions will make older ones obsolete. The rate of decay accounts for this.

Institutional changes at the patent office make it necessary to consider changes in citing behavior over time. These changes 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.

Jaffe and his co-authors develop a model that estimates the probability of a patent granted in a given year being cited by future patents, taking into account the three factors mentioned above. This paper extends their model by adding the parameter to control for technology group size. I estimate the probability of citation among patents in a cohort, p, as a function of several factors:

- the usefulness of the knowledge represented in the patent being cited ($\alpha_{i,CTD}$),
- the frequency by which patents applied for in the citing year cite earlier patents (α_{CTG})[11],
- the frequency of citations within each technology group (γ_i),
- the rate at which the knowledge represented by the cited patent becomes obsolete (β_1),
- the rate at which this knowledge diffuses through society (β_2).

Note that the first parameter, $\alpha_{i,CTD}$ is the value of most interest for this paper. It is $\alpha_{i,CTD}$ that tells us the likelihood that patents from year *CTD* will be cited by subsequent patents. The other 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. This implies that the knowledge embodied in those patents is particularly useful. Referring back to Figure 4, patents cited more frequently are at the left of the distribution of potential research outcomes. 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 $\alpha_{i,CTD}$ can be interpreted as a measure of the expected return of R&D. I refer to $\alpha_{i,CTD}$ as the *productivity parameter*.

A. Estimating the Model

The data set to be used for estimation are the citations in the energy patent data set first described in Popp (1997, 1998). It includes patents granted in the United States in 12 technology groups related to energy efficiency. The technology groups are created from various subclasses of the U.S. patent classification system. Patents are sorted by the year in which they were applied for. Although that data set includes patents applied for since 1970, only patents granted since

1975 have citation information available. A more thorough description of the data set can be found in appendix B.

Figure 5 provides examples of the trends in citation data. The graphs plot the probability of patents granted in year *CTD* being cited by other patents *x*-years after the year of grant for solar energy and fuel cell patents. In these graphs, only citations to patents in the same technology group are considered. Similar plots are obtained when all citations are used. 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, $\alpha_{i,CTD}$, 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. The figure shows 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.[12] Finally, 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.

Estimation of the values of the productivity parameter, $\alpha_{i,CTD}$, proceeds as follows. As in the work of Jaffe *et al.*, an exponential distribution is used to capture the probability of citation among cohort members. Adding an error term, $\varepsilon_{i,CTD,CTG}$, the probability of citation among cohort members can be written as:

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

Equation (3) 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.[13] To identify the parameters, it is necessary to normalize one of each of the α's to be 1. For $\alpha_{i,CTD}$, patents granted in 1970 are normalized to 1. Thus, estimates greater than one mean that patents granted in those years were more useful to future inventors than patents from the base year of 1970. Finally, since this is grouped data, the observations are weighted by $(n_{CTD}*n_{CTG})^{0.5}$ to avoid problems with heteroskedasticity (Greene, 1993).

Two separate regressions were done. In the first, only citations to other patents in the technology group are considered. Theoretically, the setup for this equation is preferable, since these are the patents represented in the distribution of potential outcomes described in section II. However, the usefulness of knowledge from other technology classes is also of interest to us. For example, advances in chemistry could lead to recharge in many energy fields. Thus, a second estimation is done in which citations to all patents are considered.

B. Results

We are most interested in the results for the productivity parameters, α_{CTD}. A high α_{CTD} means that a patent from year *CTD* is more likely to be cited by future patents, implying 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. Returning back to figure 4, a patent which is cited often, such as K_2, must be further left in the distribution of possible outcomes than a patent which is cited less often, such as K_1, leaving a greater probability for future research success. Researchers following up on the invention represented by K_2 face a higher expected marginal return to their research investment. Thus, lagged values of α_{CTD} can be interpreted as a measure of the expected return of research.

The results of the estimation using only citations within the technology groups are presented first. Figure 6 displays the results. The productivity parameters for each technology group are plotted. Numerical results for the remaining parameters are presented after the plots.

We begin with a brief look at the productivity parameter estimates. A more detailed analysis is put off until section V. Recall that the productivity parameters for 1970 were normalized to 1. Estimates greater than one for a given year indicate that patents granted in that year are more likely to be cited than patents that were granted in 1970. Note that a downward trend in the productivity parameter occurs in most groups. *The downward trend suggests that there are diminishing returns to energy research over time. It is consistent with the notion that invention potential exhaustion occurs within the individual technology groups.* There are a couple of interesting patterns. Note, for example, the peak in research opportunities before the energy crisis in solar energy. The combination of technological opportunity and higher prices together apparently account for the particularly strong reaction of solar patents to the energy crisis. Returning to figure 5, we see that patents granted in 1975 are more likely to be cited than patents from other vintages, which confirms the estimated productivity parameter.

Periods of recharge are identified by jumps in the estimates. For example, the productivity estimates of many of the technology groups exhibit an upswing at the end of the 1980's.[14] As will be discussed more thoroughly in section V, the upswing in the productivity parameters corresponds with a general upswing in the ratio of patents to R&D. Also of interest is evidence of recharge in specific technologies, such as in fuel cells during the late 1970's. Fuel cell patenting activity was strong in the 1960's, slowed during the 1970's, and then rose again in the 1980's. These results suggest that one reason for the decline in fuel cell patents in the 1970's is that research opportunities had been exhausted. New opportunities presented themselves at the end of the decade led to an increase in inventive activity in the 1980's.

Finally, in the interest of conciseness, standard errors were not presented for all of the productivity parameters. However, it should be noted that the estimates were fairly precise. In general, standard errors are about one-third to one-fourth of the estimated value. T-tests that the parameters do not equal one (the test for insignificance in this case) reveal that most estimates significantly differ from one soon after moving away from 1970.

Other results of the estimation are also encouraging. The adjusted R-square is 0.778. Both the rates of decay and diffusion are similar to the estimates in Jaffe and Trajtenberg (1995). As expected, the estimated γ's are higher for smaller technology groups. In addition, there is generally a rising trend in the citing behavior over time, α_{CTG}. However, the initial dip in α_{CTG} is unexpected. This is likely the result of "demand-side" influences in these time periods. It suggests that the decay rate is underestimated. The intuition is as follows. In the years with low α_{CTG} there are many more patents than other years. The estimated decay rate is for time elapsed. However, it is not merely time that causes the knowledge contained in a patent to lose value. Rather, patents become less valuable as they are replaced by new patents that make them obsolete. This will occur more quickly if more innovations are occurring.[15] Since the estimated decay rate is too low for years with many patents, the estimated probability of citation is too high. The estimated α_{CTG} is lower in these years to compensate.

I now move to the results of the second regression. In this regression, citations to all patents, both in and out of the technology group, were considered. Here, the scaling factor, γ, has a slightly different interpretation. In this case, it is not realistic to consider n_{CTD}, the number of potentially cited patents, to be all patents granted in a given year. Most patents are not related. Energy patents will not cite a shoelace patent, for example. To account for this, the number of patents in a given technology group is used as a measure of the potentially cited patents. This is not a perfect measure, since obviously some patents outside of the group will be cited. For example, solar energy patents might cite patents relating to semiconductors. However, it seems reasonable to assume that the pool of potentially cited patents will move in proportion to the number of patents in each energy group.

Once again, however, because of the variation in group sizes, the proportion will vary across groups. Denoting the actual pool of citable patents as $n^*_{i,CTD}$, we are claiming that $n^*_{i,CTD} = \gamma_i n_{i,CTD}$. Again, we expect γ_i to be lower for larger technology groups, since those groups would include more of the cited patents in them. For example, the solar energy technology group includes many different classifications, so it is likely that most of the patents cited are included in this group. In this case, n_{CTD} should be a good measure of the potentially cited patents. We would expect a low γ for solar energy. By comparison, the solar cells technology group is small. There may be many other classifications that are related to the technology, and thus cited often, but that are not energy specific. Classifications relating to general aspects of semiconductors would be an example. We would expect a high γ for solar cells. For identification, the γ for heat exchange, class 165, is normalized to 1.

Figure 7 presents the results for the second regression. The trends in the estimates are similar to the first regression. Looking at the productivity parameter estimates first, there is again some evidence of a general downward trend in the estimates, although not as pronounced as in the previous regression. In addition, U-shaped trends are more prevalent in the productivity estimates using all citations. Since the U-shaped plots are not as strong when only citations within the group are considered, the U-shaped plots may be evidence of new inventions outside of these fields providing recharge. However, other explanations are possible. Higher productivity estimates tend to coincide with a low number of citable patents. It may suggest that when R&D efforts increase, inventors are left to discover inventions at the bottom of the barrel that would have otherwise been ignored. This would suggest diminishing returns to increases in research spending *in any one time period*. It may also be the case that using only patents in the technology group as a measure of potentially cited patents might not capture all of the movement in the potentially cited patents. These explanations will be explored more fully in section V.

Other results of this regression are similar to the first. The adjusted R-square of 0.8 is slightly higher. Again, most estimates are still significantly different from one soon after moving away from 1970. One other result is of note. The decay rate is slightly smaller, 0.32 instead of 0.39. In addition, the unexpected "demand side" drop in citing practices, α_{CTG} is even larger than before. This is further evidence that the decline is due to underestimating the decay rate.

IV. Do the Productivity Estimates Measure the Expected Returns to R&D?

In this paper, the lagged values of the productivity parameters just estimated are to be interpreted as a measure of the expected returns to R&D. In this section, I test whether this conjecture is reasonable. I verify that patent citations are a valid measure of the value of a patent by examining the relationship between patent citations and patent renewals. I show that the likelihood of renewing a patent increases as the number of citations to the patent increases.

Beginning with the work of Pakes and Schankerman (1984), patent renewal data have been used to value patent protection. In most countries, patent holders must pay an annual fee to keep their patent in force. In addition, the required fee increases with age. Pakes and Schankerman note that a patent holder maximizing the expected discounted value of the net returns to patent protection will only pay the fee to renew a patent if the value of patent protection is greater than the fee. As a result, patent renewal data can be used to estimate a distribution of the value of patents. Several papers have since used this technique to value patent protection. In each, a highly skewed distribution is found, with most patents having little value, but with a few highly valuable patents. Lanjouw, Pakes, and Putnam (1996) provide a survey of the patent renewal literature.

Unfortunately, patent renewal data is of less use when analyzing patents granted in the United States. Patent renewal fees were instituted for patent applications filed in the U.S. after December 12, 1980. Until then, no renewal fee was required. In contrast, renewal fees have been required in Europe for much longer. Since the maximum length of a patent in the U.S. is 17 years, the first patents subject to patent renewal fees are just now reaching the end of their lives. Thus, useful time series can not be obtained for estimating the value of U.S. patents using renewal data. In addition, the renewal fees are only due three times during the life of a patent (after 3 1/2, 7 1/2, and 11 1/2 years after the patent was granted), so that the fees are not as large a burden as those as other countries. As a result, using U.S. patent renewal data to estimate a value of patent protection is not feasible.

However, the limited renewal data on U.S. patents does make it possible to test the validity of the productivity estimates. In particular, the productivity estimates rely on the assumption that more citations to a patent indicate that the patent is more valuable. To test this hypothesis, in this section regressions are run to compare the probability of renewing a patent to the number of citations it has received by the time renewal is required. A patent will be renewed only if the expected value of future returns is greater than the renewal fee. If citations are a good measure of the value of a patent, the number of citations should have a positive effect on the probability of renewal.

For each technology group, a logit regression is run relating the probability of renewal in the eight year to the number of citations received by the eighth year.[16] The dependent variable, y, is equal to one if the patent was renewed after eight years, and 0 if it was allowed to expire. The independent variable is the number of citations received by the patent in the first eight years of its life. The function estimated is:

$$(4) \qquad E[y] = \Lambda(\boldsymbol{\beta}'\mathbf{x}),$$

where:

$$y = \begin{cases} 1 \text{ if the patent was renewed in the 8th year} \\ 0 \text{ otherwise} \end{cases}$$

$$\mathbf{x} = [1, \# \text{ of citations received after 8 years}]'$$

$$\Lambda = \frac{e^{(\boldsymbol{\beta}'\mathbf{x})}}{1+e^{(\boldsymbol{\beta}'\mathbf{x})}} = \text{the logistic cumulative distribution function}$$

Table 1 presents the marginal effects of an additional citation on the probability of renewal. The effect is positive for each technology group. The standard errors of the estimations are low, so that most coëfficients are significant. Seven of the twelve are significant at the one percent level, and all but two are significant at the five percent level. The range of the marginal effects is between 1.5% and 12.5%. Both the mean and median increase is approximately four percent.

These results suggest that there is a significant correlation between patent renewals and patent citations. Since many papers have found European patent renewal data to be a useful indicator of the value of patents, the correlation between renewals and citations suggests that patent citation data are also of use in valuing patents. This result is of particular importance to researchers using United States patent data, since patent citation data is more readily available than patent renewal data in the United States.

V. Are There Diminishing Returns to R&D?

Having demonstrated that the productivity parameters are a valid measure of the expected returns to R&D, I now turn to the final question of this paper: are there diminishing returns to R&D? For this, I will examine the behavior of the productivity estimates over time.

It is a well-documented fact that the ratio of patents to real R&D expenditures declined in the industrialized world during the 1970's and 1980's, before turning upward again in the late 1980's.[17] Figure 1 illustrates this decline. Three main explanations have been offered for this decline. Based on the search model of R&D, Evenson (1991) offers *invention potential exhaustion* (IPE) as a possible explanation. As discussed in section II, the IPE explanation observes that inventors have a limited pool of possible inventions from which new innovations are drawn. As more inventions are created, fewer possibilities for future success exist. The search for new ideas gets harder and harder. Griliches (1989, 1990) disputes this view, arguing instead that changes in the propensity to patent inventions have caused the decline in patent-to-R&D ratios. According to this argument, it is the likelihood of a successful inventor seeking patent protection that has fallen, not the likelihood of successful research. Finally, Kortum (1993) looks at the effect of increased demand on returns to R&D. If demand for research increases, marginal projects, which did not seem worthwhile before, will be undertaken. The overall productivity of research will fall as a result.

The downward-sloping plots of the productivity estimates suggest that there are diminishing returns to R&D. In terms of the search model of R&D presented in section II, the estimates suggest that researchers are moving to the right of the distribution as time progresses, and that new knowledge is not causing the distribution to shift fast enough to compensate. In this section, I take a closer look at the productivity estimates and the hypothesis of diminishing returns. First, I regress the productivity estimates, $\alpha_{i,CTD}$, on time and the total number of patents. Next, a quick comparison of the productivity estimates to the ratio of patents to R&D spending in energy technologies is presented. Finally, the implication of these results on the patent-to-R&D expenditure ratio is discussed.

A. Regression Analysis

Diminishing returns to the productivity of research could exist in one of two forms. First, there may be diminishing returns over time. As more inventions are created, the pool of potential new inventions dries up, making future successes more difficult. Evenson's invention potential exhaustion hypothesis predicts this.[18] Second, the returns to additional R&D spending in any given period may be diminishing. Intuitively, since optimizing researchers will choose the most productive avenues of study first, additional research must focus on less and less productive searches. Formally, the implication is that the second derivative of the knowledge with respect to R&D, $\partial^2 K_t / \partial R^2_t$, is negative. This section searches for evidence of each.

For this, two additional regressions were run. First, the productivity estimates were simply regressed on a constant and a time index. A negative coefficient on the time index is evidence of diminishing returns over time. Second, the productivity estimates were regressed on a constant, a time index, and the number of patents granted in each year. A negative coefficient on the number of patents is evidence of diminishing returns to additional research spending within a given year. The data from each technology group were pooled so that single coefficient estimates could be obtained. The results were corrected for autocorrelation.

I begin by testing for diminishing returns across time. For this, the equation estimated is:

(5) $$\alpha_{i,CTD} = a + b \cdot t + \varepsilon_{it}, \quad i = 1,\ldots,N;\ t = 1,\ldots,20.$$

A significantly negative estimate of b is evidence that the productivity of R&D is falling over time. The productivity estimates represent how useful the ideas contained in patents of a given year are to future inventors. A declining productivity estimate implies that the ideas are less useful, meaning that researchers are moving rightward in the distribution of possible inventions.

Results are presented in the first two columns of table 2. Column 1 presents the results for the productivity estimates that were obtained by using citations to other patents in the same technology group, and column 2 uses the productivity parameters obtained by including all patent citations. For the productivity estimates using only within group citations, the sign of b is negative and significant at the 10 percent level. *This is evidence that diminishing returns over time – or invention potential exhaustion – occurs on research within individual technology groups.*

For the regression using estimates of citations to all patents, presented in the second column, opposite results are found. The coefficient on the time trend is positive, although it is not significantly different from zero. Recall that many of the productivity estimates using all patents exhibited U-shaped trends. *Thus, there is no evidence of diminishing returns across time when looking at the usefulness of all patents.*

Finally, a second set of regressions was run using both time and the number of patents in a given year as dependent variables. The number of patents is included to test for diminishing returns to R&D *within* a given year. If the coefficient on the number of patents is significant and negative, it reveals that the average patent is less likely to be cited in years when more inventions occur. This suggests that some of the innovations are from the "bottom of the barrel" of new ideas, or that research is being duplicated, so that there are more similar patents that are potential cites for future patents. Observing that R&D expenditures have increased, rather than decreased, over time, Kortum (1993) suggests that diminishing returns to R&D in a given year, rather than across time, may explain the falling patent-to-R&D ratio. A negative coefficient on the number of patents would be consistent with Kortum's hypothesis. The estimated equation is:

(6) $$\alpha_{i,CTD} = a + b \cdot t + c \cdot P_t + \varepsilon_{it}, \quad i = 1,\ldots,N;\ t = 1,\ldots,20,$$

where P_t represents the number of successful patent applications in year t.

Results are presented in the last two columns of table 2. Again, the results for citations within the technology group are presented first. Including the number of patents in a given year offers a slight improvement to the adjusted R-square. Again, the coefficient on time, b, is negative and significant. For citations within the technology group, the coefficient on the number of patents, c, is positive and highly significant.[19] These results suggests that current research does use up ideas for future researchers, but that an increase in current R&D is not likely to lead to redundant research. Of course, this result could be clouded by the nature of the patenting process. It may be that redundant research is being carried out, but that only one patent is being awarded for each similar invention.

For the set of regressions using all citations, the results are opposite. The time trend, b, is positive. However, the number of patents, c, has a negative impact, and is highly significant. In addition, there is a dramatic improvement in the adjusted R-square. For this regression, the productivity estimate $\alpha_{i,CTD}$ represents the productivity of *all* patents granted in year *CTD*. These results suggest that, when considering the usefulness of all technologies, diminishing returns to research occur within a given year, rather than over time. It appears that recharge occurs to keep the returns to research from falling too drastically over time.

Thus, the results suggest a bit of a contradiction. To summarize, I find:

1) evidence of diminishing returns over time when looking only at patent citations within a technology group, with little evidence of diminishing returns to additional research in a given year, but
2) evidence of diminishing returns to additional research in a given year when looking at citations to all patents, with little evidence of diminishing returns to research over time.

How can this apparent contradiction be explained? An important factor is likely to be the ability of scientific advances outside the technology fields to keep pace with changes in the level of innovative activity within these fields. The number of patents in the technology groups is strongly influenced by demand forces such as energy prices. However, higher energy prices should not lead to an increase in useful knowledge from outside the technology field. As a result, researchers in the energy fields will be trying to get more ideas out of a fixed base of outside knowledge. An increase in innovation in any period will cause this knowledge to be used up more quickly, leading to diminishing returns in any given time period when looking at the usefulness of outside knowledge.

In contrast, the observation that the number of patents granted does not have much affect on the likelihood of patents within the technology group being cited suggests that demand forces that encourage innovation prevent technology-specific knowledge from being used up too quickly. New advances within the field are created more quickly when energy prices are high. However, these appear to be useful innovations for future inventions. The "bottom of the barrel" is not reached quickly. It is reached eventually, however, as each year slightly less valuable ideas are created. This may suggest that the knowledge becomes more specialized as time progresses, so that it is not as useful to future inventors. Finally, the lack of evidence for recharge in the within group citations suggests that the new inventions that result from recharge outside of the technology group are not as useful as previous inventions.

B. Energy Patents and R&D Spending

The energy patent data set also offers insight into the falling patent-to-R&D ratio, at least in the narrowly defined field of energy research. For this, a measure of energy R&D expenditures to use with the data on patent counts is needed. Unfortunately, published R&D data do not go into great detail as to the nature of the research. However, the National Science Foundation does provide estimates of the amount of industrial R&D expenditures devoted to energy research. The data are taken from the NSF's Survey of Industrial Research and Development. Included in the survey results are the amounts of energy R&D expenditure by primary energy source. These figures are presented in table 3.

In this paper, I have argued that lagged estimates of $\alpha_{i,CTD}$ provide information on the productivity of R&D inputs. I will use these estimates in combination with the ratio of patents to R&D expenditures in selected technologies to see what decline, if any, in the ratio can be explained by the productivity estimates. Figure 8 plots the productivity estimates from section III, lagged one year, along with the patent-to-R&D ratios in three fields: solar energy, coal, and various demand-side technologies.[20]

The figures show that changes in productivity explain much of the variation in these ratios. I have asserted that citations to patents granted in a given year are an indicator of the remaining possible inventions. Thus, years with higher estimated $\alpha_{i,CTD}$'s should precede years in which the ratio of patents to R&D expenditures is high. In general, this appears to be the case. Several features stand out. First are the high patent-to-R&D ratios for solar energy during the

1970's, when research on these technologies was in its infancy. The productivity estimates confirm that research was productive in this period.

Second, note that the patent-to-R&D ratio rises in the late 1980's for the conservation and coal technologies. In addition, note that solar energy does not experience a rise in either the productivity estimates or the patent-to-R&D ratios during the late 1980's. These results seem to confirm the U-shaped plots of the productivity estimates that are found in many of the technologies. Two possible explanations exist. One is that the U-shaped plots are evidence of recharge. New inventions from outside these fields may have made research more productive. Interestingly, a similar pattern is observed in the ratio of patents to R&D for all patents and R&D expenditures in the U.S. during this period. Another possibility is that demand forces are at work. Note the low patent-to-R&D ratios for conservation technologies in the 1970's. Whereas the supply side technologies, such as solar energy, were in their infancy at this time, the conservation technologies did exist before the energy crisis. The energy crisis merely stimulated research activity in these fields. Thus, while researchers on solar energy were dealing with an untapped pool of new ideas, researchers in the conservation field may have been faced with diminishing returns. By increasing their R&D investments in any given year, researchers were left with ideas at the "bottom of the barrel". These ideas, being less valuable, are cited less often, resulting in lower productivity estimates.

C. Implications for the Patent-to-R&D ratio

The above results shed much light on the puzzle of the falling patent-to-R&D ratio. At first glance, diminishing returns, as a result of exhaustion, would appear to be the culprit. Although this explanation seems satisfactory to explain behavior in the energy innovation field, it does not seem enough to explain the fall in the overall ratio. As noted by Griliches (1989):

> A priori one would expect to hit diminishing returns in any narrowly defined field, at least until the field or the product area is redefined anew by another major breakthrough.... Inventive effort, however, moves from one fishing ground to another, and new fishing grounds open up as the result of basic R&D and other sources of discovery. Thus, in the longer run, there is less evidence of exhaustion of opportunities, and studies that have tried to look for declines in the rates of return to R&D have found very little evidence of them. (p. 317)

Griliches argument implies that falling productivity estimates will result in R&D shifting to more productive sectors. Popp (1998) confirms this by showing that higher lagged values of the productivity estimates lead to more patenting activity in the energy technology fields studied in this paper. Since research inputs flow towards more productive sectors, evidence of diminishing returns within a technology field is not enough to prove that the falling patent-to-R&D ratio observed in the entire economy is also due to diminishing returns. For this to be true, we would also need evidence that new sectors of research are less productive than older ones. Other studies such as this one could be useful for that purpose.

VI. Conclusion

This paper used patent citation data to examine trends in the productivity of energy research from 1970-1990. It estimated annual productivity parameters for patents in 12 energy technology groups. The major findings of this work are:

- The likelihood of subsequent citations to a newly granted patent, as captured by the productivity estimates in section III, falls over time.
- The productivity estimates are correlated with changes in the ratio of patents to R&D spending, suggesting that they are a useful measure of the productivity of R&D spending.
- Declining productivity estimates are suggestive of invention potential exhaustion within each technology group.

- In addition, after a spurt of innovative activity, the sciences outside of a technology group need time to catch up. Thus, some recharge occurs after an increase in innovative activity. However, the new inventions that follow the recharge tend to be less useful than older inventions.

The lessons drawn from this paper have implications both for environmental economists and technology economists. For environmental economists, evidence of diminishing returns to R&D suggest that the largest stimulus to innovation after the imposition of a carbon tax or other similar policy will come soon after the policy is implemented. For the stimulus to innovation to continue, research opportunities must be maintained. Government-supported R&D, by increasing the stock of basic knowledge, could play a role in this.

For technology economists, the results shine some light on the mystery of declining patent-to-R&D ratios in the industrial world. We see that the productivity estimates in the energy technologies decline during the time frame studied in this paper, and that this decline seems to correspond to falls in the patent-to-R&D ratio. However, this alone cannot explain the falling patent-to-R&D ratio. Rational investor behavior implies that less innovative activity will occur in fields with low productivity estimates. Inventive activity should shift toward more productive areas of research. This is confirmed in Popp (1998), which finds that lagged values of the productivity estimates are good predictors of innovative activity. For diminishing returns to lead to a falling patent-to-R&D ratio, it would have to be the case that new areas are less productive than old areas of. Micro-level studies of other technologies would be useful to resolve this question.

Finally, the results of this paper can also be linked to the theory of endogenous growth. Endogenous growth depends on the returns to research not diminishing, so that incentives to improve knowledge always exist. Having shown that diminishing returns exist within narrowly defined technology fields, knowing how the returns to research vary across different areas would be of use to endogenous growth theory as well.

APPENDIX A – THE SEARCH MODEL OF R&D

The search model of R&D, first proposed by Evenson and Kislev (1975), was introduced in section II. In this appendix, a more formal mathematical treatment of the model is presented. The model treats research and development as a search process over a random distribution of outcomes. This distribution is unknown to us, but we assume that a firm undergoing research has an idea of what the probability for success is before beginning a research project. Denote this distribution of possible outcomes as $f(x|\mu,\sigma)$, a probability density function of x with mean μ and variance σ^2. The current state of the technology, which is the highest outcome from the previous draws, can be denoted K_t.

In the model, R&D results in random draws being taken from this distribution. Research is only successful if the result is greater than the current state of technology. At any time t, n_t searches are done over this distribution. For this paper, n can be thought of as the amount spent on research and development activity. x_i is the yield of the ith search. The research is successful only if the result is higher than the previous best state of the technology, or if $x_i > K_t$. For simplicity, assume that a firm only produces one innovation per period, so that for a given set of searches with results x_i, where $i = 1, 2,\dots,n$, the only result in which the firm is interested in is the highest. Denote this as z, so that $z = x_j$, $x_j \geq x_i \ \forall \ i$. Let $F(x|\mu,\sigma)$ denote the cumulative distribution of $f(x|\mu,\sigma)$. Using the theory of order statistics, the cumulative distribution of z is

(A1) $$J_n(z|\, \mu,\sigma) = \Pr\,(\text{all } x_i \leq z|\, \mu,\sigma) = F^n(z|\, \mu,\sigma)$$

and the probability density function is

(A2) $$j_n(z) = n\, F^{n\text{-}1}(z|\, \mu,\sigma)f(z|\, \mu,\sigma).$$

Note that the probability density function is, among other things, a function of the number of draws taken. Additional draws in a given period shift the distribution of potential outcomes outward. See figure A1.

To value innovation, normalize K and z so that the distance between them is the value of the innovation. This gives us:

$$\text{(A3)}\; EV = \varphi(R_t, K_t) = E[K_{t+1} - K_t] = \int_{K_t}^{\infty} z j(z|\mu_t, \sigma_t)dz = \int_{K_t}^{\infty} z R_t F^{R-1}(z|\mu_t, \sigma_t) f(z|\mu_t, \sigma_t)dz,$$

where R_t, the amount spent on R&D at time t, replaces n. The expected value of research is illustrated in figure 2. The shaded area is the expected value of research given that the current level of technology is at K_t. It represents the possible successes that have not yet been realized. We can think of this area as the stock of unexploited technology.

Because I do not observe the distributions in equation (A3), I cannot directly calculate values for the expected value of research at time t. For this, data on both research inputs and outputs would be needed. However, as discussed in section II, it is possible to use the patent citations to observe how the expected value changes over time. I review that discussion here, to make clear how the use of citations is linked to the search model of R&D. Consider a distribution of potential outcomes for a given number of draws, N^*. With a current level of technology of K_t, the expected value of research is the shaded area in figure 2. Now, consider what happens in the next period. If a successful invention occurs in period t, the state of knowledge in period t+1 shifts rightward. If there is no change in the distribution of potential outcomes, the expected value of research will be less in period t+1. This is illustrated in figure 3a. The expected value of research in time t is equal to A+B. In period t+1, after the level of knowledge has shifted outward, the expected value is only area B. Evenson (1991) refers to this as *invention potential exhaustion.* Formally, this relationship is derived from equation (A3) as:

(A4) $$\varphi_K = -zR_t F^{R-1}(K_t|\mu_t,\sigma_t) f(K_t|\mu_t,\sigma_t) < 0.$$

As a result, the marginal returns to R&D are lower after the shift of knowledge:

(A5) $$\varphi_{RK} = -zF^{R-1}(K_t|\mu_t,\sigma_t) f(K_t|\mu_t,\sigma_t) - zR(R-1)F^{R-1}(K_t|\mu_t,\sigma_t) f(K_t|\mu_t,\sigma_t) < 0.$$

However, it is possible that new knowledge will change the nature of the distribution. New innovations may contain basic knowledge, which make possible in period t+1 inventions that were not possible in period t. Evenson refers to this as *recharge*. He notes that recharge can come from inventions in related fields or from advances in basic science. For example, the invention of microcomputer chips led to the development of a generation of electronic equipment. In the search model, recharge is represented as an increase in the mean of the distribution. Formally, $\mu_{t+1} > \mu_t$. Another possibility is that the results of innovation may make clearer which avenues of research are likely to be successful and which aren't. Thus, future researchers will have a better idea of which research to pursue, lowering the variance of future results. Formally, $\sigma_{t+1} < \sigma_t$. In this case, even unsuccessful research can be helpful, by clarifying which direction future research should take.

Recharge is illustrated in figure 3b. Again, the original state of knowledge is K_t, and the new state of knowledge is K_{t+1}. However, in this example, the distribution of potential outcomes has shifted outward. The expected value of research in period t is A+B; in period t+1 it is B+C. Since area C is drawn to be larger than area A, the effect of recharge dominates the exhaustion in this example. As Evenson notes, the debate over returns to R&D centers on the possibilities for recharge. The principles of diminishing returns should apply to research just as much as any other economic activity, leading to invention potential exhaustion. The question is whether new advances in science occur quickly enough to counteract the effect of exhaustion. In this paper, instances in which the effect of recharge dominates invention potential exhaustion are observed as positive jumps in the estimates of the productivity parameters.

APPENDIX B – DATA DESCRIPTION

The most important data used in this paper is the energy patent data set. To create the energy patent data set, I first created 12 technology groups pertaining to energy efficiency, with each group containing several U.S. patent classifications. Appendix C briefly describes the major technology areas studied, and Appendix D lists the technology groups and the patent classifications used in each. The data set includes all successful patent applications in these classifications in the United States from 1970 to 1994. The patent data set was assembled from two sources. The main source of patent data for this project is the MicroPatent CD-ROM database of patent abstracts. The MicroPatent database contains every U.S. patent issued from 1975-1994. This data set includes all of the information that is found on the front page of a patent, including the date of application, date of grant, name of the inventor and his or her organization, the nation of the patent holder, and citations to previous patents. In addition, since the MicroPatent database does not extend far enough into the pass to encompass the first energy crisis of 1973, additional lists of patents in our target classes has been obtained from the Classification and Search Support System (CASSIS) available from the U.S. Patent and Trademark Office. Unfortunately, for these patents the additional data found on the MicroPatent database is not available.

Using these data sources, all patents in the classifications in question were chosen. Both patents granted to American and foreign inventors are included in this paper, since the origin of knowledge should not affect its usefulness much. For each technology group, patents are sorted by the year of application. Several papers have found that patents, grouped by the date of application, to be a good indicator of R&D activity (see Griliches 1990 for a survey). Figure B1 shows trends in the patent data from 1969 to 1993, along with an index of energy prices taken from the *State Energy Price and Expenditure Report*. Since a patent application is only made public when a patent is granted, the data for the last 10 years have been scaled up by using a distribution of the lag between patent application dates and patent grant dates for the patents in the sample. Note that for most of the technology groups, there is a jump in patent applications during the energy crises of the 1970's. Notable exceptions include fuel cells, the use of waste as fuel, and continuous casting. Reasons why these technologies differ are discussed in Popp (1998).

APPENDIX C – ENERGY TECHNOLOGY GROUPS STUDIED

For this study, 12 groups of energy-related patents were created, with each group representing a different technology. Of these, 7 are related to energy supply, and 5 are related to energy demand. Groups were included based on their importance to energy conservation and the clarity of patent classification. Importance to energy conservation was determined by a review of literature on energy conservation, including references on industrial energy conservation and several Department of Energy publications. Clarity of patent classification means that the patent classifications included must clearly be used for energy conservation purposes.[21] Appendix D lists the patent classifications in each of the 12 technology groups.

Supply technologies included in the data set were those that are substitutes for fossil fuels. Nuclear technologies were not included, since nuclear patents would be strongly influenced by changing attitudes about the safety of nuclear power and the resulting changes in the regulatory burden of the nuclear industry. Technologies that were utilized include technologies for getting fuel from coal, fuel cells, and renewable sources such as solar, wind, and biomass energy.

Demand technologies focus mainly on industrial energy consumption. They include a mixture of technologies chosen for specific industries, and technologies with a more general usage. Energy use varies greatly across sectors. Petroleum refining, chemicals, primary metals, pulp and paper, food, and ceramics and glass accounted for 74% of total industrial energy use in 1988. In some cases, energy use is high because the output of the industry is high – food and oil and gas extraction are examples. In other industries, both output and the energy intensity is high. Such industries are petroleum refining, steel, organic chemicals, and paper (OTA, 1993). Table C1 presents energy consumption and energy intensity in selected industries. Energy intensity is calculated as real energy expenditures divided by real value added in the industry. Most energy R&D is concentrated in these energy intensive industries.

Table C1 - Energy Use for Selected Industries

Industry	SIC	Real Energy Consumption (millions of 1987 $'s)			Energy Intensity (Real Energy/Real Value Added)		
		1971	1981	1991	1971	1981	1991
Food	20	3,916	4,534	4,597	0.05	0.05	0.03
Paper	26	4,289	4,970	5,184	0.13	0.13	0.10
Chemicals	28	8,028	10,389	8,174	0.10	.012	0.06
Petroleum Refining	29	3,836	4,214	4,176	0.17	0.26	0.18
Stone, Clay, and Glass	32	3,932	4,027	3,194	0.13	0.14	0.10
Primary Metals	33	10,225	10,765	7,008	0.18	0.21	0.17
Steel	3312	6,064	5,741	2,951	0.22	0.27	0.21
Aluminum	3334	1,101	1,539	1,266	0.53	0.65	0.70

Source: Calculated from data in NBER Manufacturing Productivity Database

Of the demand-side technology groups, some have general industrial applications. Industrial processes produce much heat that is never used. Recovery of this waste heat would allow for substantial energy savings. Three of the categories relate to the use of waste heat – waste heat, heat exchange, and heat pumps. Some of the technology groups chosen relate to specific energy intensive industries. For example, continuous casting is a method of steel production that requires less energy than traditional techniques. Industry-specific technologies could not be chosen for all of the energy intensive industries, because in some of these industries there are a broad range of energy uses, making it difficult to focus on individual technologies. For example, the chemical industry makes a diverse range of products such as rubber, plastics, soaps, paints, industrial gases, fertilizers, and pharmaceuticals. Obviously, different technologies to conserve energy would be needed in the production of these various products. In addition, for products such as these, it is difficult to identify innovations that are clearly related to energy conservation, as opposed to other improvements in the product.

APPENDIX D: U.S. patent classifications related to energy

Guide to definitions: The first phrase is the main classification. For example, class 208 contains patents for Mineral Oils: Processes and Products. These are followed by the various subclassifications, listed in descending order of precedence.

Supply Technologies:

Coal Liquefaction:

208/400-435	Mineral Oils: Processes and Products/By treatment of solid material (e.g. coal liquefaction)

Coal Gasification:

48/200	Gas: Heating and Illuminating/Processes/Coal, oil and water
48/201	Gas: Heating and Illuminating/Processes/Coal and oil
48/202	Gas: Heating and Illuminating/Processes/Coal and water
48/210	Gas: Heating and Illuminating/Processes/Coal
48/71	Gas: Heating and Illuminating/Generators/Cupola/Coal, oil and water
48/72	Gas: Heating and Illuminating/Generators/Cupola/Coal and oil
48/73	Gas: Heating and Illuminating/Generators/Cupola/Coal and water
48/77	Gas: Heating and Illuminating/Generators/Cupola/Producers/Coal
48/98	Gas: Heating and Illuminating/Generators/Retort/Coal, oil and water
48/99	Gas: Heating and Illuminating/Generators/Retort/Coal and water
48/100	Gas: Heating and Illuminating/Generators/Retort/Coal and oil
48/101	Gas: Heating and Illuminating/Generators/Retort/Coal

Solar Energy:

60/641.8-641.15	Power Plants/Utilizing natural heat/Solar
62/235.1	Refrigeration/Utilizing solar energy
126/561-568	Stoves and Furnaces/Solar heat collector for pond or pool
126/569-713	Stoves and Furnaces/Solar heat collector
126/903	Stoves and Furnaces/Cross-Reference Art/Solar collector cleaning device
126/904	Stoves and Furnaces/Cross-Reference Art/Arrangements for sealing solar collector
126/905	Stoves and Furnaces/Cross-Reference Art/Preventing condensing of moisture in solar collector
126/906	Stoves and Furnaces/Cross-Reference Art/Connecting plural solar collectors in a circuit
126/910	Stoves and Furnaces/Cross-Reference Art/Heat storage liquid

Solar Energy – Batteries:

136/206	Batteries: Thermoelectric and Photoelectric/Thermoelectric/Electric power generator/ Solar energy type
136/243	Batteries: Thermoelectric and Photoelectric/Photoelectric
136/244-251	Batteries: Thermoelectric and Photoelectric/Photoelectric/Panel
136/252-265	Batteries: Thermoelectric and Photoelectric/Photoelectric/Cells

Fuel Cells:

429/12-46	Chemistry: Electrical Current Producing Apparatus, Product, and Process/Fuel cell, subcombination thereof or method of operating

Wind:

290/44	Prime-Mover Dynamo Plants/Electric control/Fluid-current motors/Wind
290/55	Prime-Mover Dynamo Plants/Fluid-current motors/Wind
416/132B	Fluid Reaction Surfaces (i.e., Impellers)/Articulated resiliently mounted or self-shifting impeller or working member/Sectional, staged or non-rigid working member/windmills
416/196A	Fluid Reaction Surfaces (i.e., Impellers)/Lashing between working members or external bracing/Connecting adjacent work surfaces/Non-turbo machine (windmills)
416/197A	Fluid Reaction Surfaces (i.e., Impellers)/Cupped reaction surface normal to rotation plane/Air and water motors (natural fluid currents)

Using waste as fuel:

110/235-259	Furnaces/Refuse incinerator
110/346	Furnaces/Incinerating refuse

Demand Technologies:

Waste heat:

122/7R	Liquid Heaters and Vaporizers/Industrial/Waste heat
7A	Liquid Heaters and Vaporizers/Industrial/Waste heat/Steel converter
7B	Liquid Heaters and Vaporizers/Industrial/Waste heat/Additional burner
7C	Liquid Heaters and Vaporizers/Industrial/Waste heat/Waste sulfate
7D	Liquid Heaters and Vaporizers/Industrial/Waste heat/Carbon monoxide
60/597-624	Power Plants/Fluid motor means driven by waste heat or by exhaust energy from internal combustion engine

Heat exchange:

165	Heat Exchange

Heat pumps:

62/238.7	Refrigeration/Disparate apparatus utilized as heat source or absorber/With vapor compression system/Reversible, i.e. heat pump
62/324.1-325	Refrigeration/Reversible, i.e., heat pump

Stirling engine:

60/517-526	Power Plants/Motor operated by expansion and/or contraction of a unit of mass of motivating medium/Unit of mass is a gas which is heated or cooled in one of a plurality of constantly communicating expansible chambers and freely transferable therebetween

Continuous casting:

148/541	Metal Treatment/Process of modifying of maintaining internal physical structure (i.e. microstructure) or chemical properties of metal, process of reactive coating of metal and process of chemical-heat removing (e.g., flame-cutting, etc.) or burning of metal/With casting or solidifying from melt/Iron(Fe) or iron base alloy/Continuous casting
148/551	Metal Treatment/Process of modifying of maintaining internal physical structure (i.e. microstructure) or chemical properties of metal, process of reactive coating of metal and process of chemical-heat removing (e.g., flame-cutting, etc.) or burning of metal/With casting or solidifying from melt/Aluminum (Al) or aluminum base alloy/Continuous casting
164/263	Metal Founding/With product severing or trimming means/Associated with continuous casting means
164/268	Metal Founding/With coating means/associated with a continuous or semicontinuous casting means
164/415	Metal Founding/Means providing inert or reducing atmosphere/In continuous casting apparatus
164/416	Metal Founding/Including vibrator means/In continuous casting mold
164/417	Metal Founding/Combined/Including continuous casting apparatus
164/418-444	Metal Founding/Means to shape metallic material/Continuous or semicontinuous casting
164/445-446	Metal Founding/Starter bar
164/447-448	Metal Founding/Product supporting or withdrawal means for continuous casting apparatus
164/449.1-450.5	Metal Founding/Control means responsive to or actuated by means sensing or measuring a condition or variable (i.e., automatic control)/Control of feed material enroute to shaping area/Responsive to material level/In continuous casting apparatus
164/451-455	Metal Founding/Process/With measuring, testing, inspecting, or condition determination/Of continuous or semicontinuous casting
164/459-491	Metal Founding/Process/Shaping liquid metal against a forming surface/Continuous or semicontinuous casting
164/502-504	Metal Founding/Including means to directly apply magnetic force to work or to manipulate or hold shaping means/In continuous casting apparatus
164/505-509	Metal Founding/Means to directly apply electrical or wave energy to work/In continuous casting apparatus
164/154.4	Metal Founding/Control means responsive to or actuated by means sensing or measuring a condition or variable (i.e., automatic control)/Responsive to position or spatial dimension/Responsive to rate of change/Continuous casting
164/154.5	Metal Founding/Control means responsive to or actuated by means sensing or measuring a condition or variable (i.e., automatic control)/Responsive to position or spatial dimension/Continuous casting

REFERENCES

Binswanger, Hans P. (1978), "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.

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.

Clark, Kim B. and Zvi Griliches (1984), "Productivity growth and R&D at the business level: results from the PIMS data base," in Zvi Griliches, ed., *R&D, Patents and Productivity*, University of Chicago Press, Chicago, 375-92.

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

Evenson, Robert E. and Yoav Kislev (1975), *Agricultural Research and Productivity*, Yale University Press, New Haven, CT.

Greene, William H. (1993), *Econometric Analysis*, Macmillan Publishing Co., New York.

Griliches, Zvi (1995), "R&D and Productivity: Econometric Results and Measurement Issues," in Paul Stoneman, ed., *Handbook of the economics of innovation and technological change*, Blackwell Publishers, Cambridge, MA, 52-89.

Griliches, Zvi (1990), "Patent statistics as economic indicators," *Journal of Economic Literature*, 28, 1661-1707.

Griliches, Zvi (1989), "Patents: Recent Trends and Puzzles," *Brookings Papers on Economic Activity: Microeconomics*, vol. x, 291-330.

Griliches, Zvi and Jacques Mairesse (1984), "Productivity and R&D at the firm level," in Zvi Griliches, ed., *R&D, Patents and Productivity*, University of Chicago Press, Chicago, 339-74.

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., 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 Manuel Trajtenberg (1995) "Flows of Knowledge From Universities and Federal Labs: Modeling the Flow of Patent Citations Over Time and Across Institutional and Geographic Boundaries," National Academy of Sciences Colloquium on Science, Technology and the Economy.

Jung, Chulho and Krutilla, Kerry (1996), "Incentives for Advanced Pollution Abatement Technology at the Industry Level: An Evaluation of Policy Alternatives," *Journal of Environmental Economics and Management*, 30, 95-111.

Kortum, Samuel (1993), "Equilibrium R&D and the Patent-R&D Ratio: U.S. Evidence," *American Economic Review*, 83, 450-457.

Lanjouw, Jean O. (1993), "Patent Protection: Of What Value and For How Long?" NBER Working Paper no. 4475.

Lanjouw, Jean O., Ariel Pakes, and Jonathan Putnam (1996), "How to Count Patents and Value Intellectual Property: Uses of Patent Renewal Data and Application Data," NBER Working Paper no. 5741.

Milliman, Scott R. and Raymond Prince (1989), "Firm Incentives to Promote Technological Change in Pollution Control," *Journal of Environmental Economics and Management*, 17, 247-265.

Milliman, Scott R. and Raymond Prince (1991), "Firm Incentives to Promote Technological Change in Pollution Control: Reply," *Journal of Environmental Economics and Management*, 22, 292-296.

NBER Manufacturing Productivity Database, available on-line.

Nelson, Richard R. and Sidney G. Winter (1982), *An Evolutionary Theory of Economic Change*, The Belknap Press of Harvard University Press, Cambridge, MA.

Newell, Richard G., Adam B. Jaffe, and Robert N. Stavins (1996), "The Induced Innovation Hypothesis and Energy-Saving Technological Change," NBER working paper #6437.

Office of Technology Assessment (1993), *Industrial Energy Efficiency*, Washington, D.C., U.S. Government Printing Office, August.

Pakes, Ariel (1986), "Patents as options: some estimates of the value of holding European patent stocks," *Econometrica*, 54, 755-84.

Pakes, Ariel and Mark Schankerman (1984), "The rate of obsolescence of knowledge, research gestation lags, and the private rate of return of research resources," in Zvi Griliches, ed., *R&D, Patents and Productivity*, University of Chicago Press, Chicago, 73-88.

Popp, David (1998), "Induced Innovation and Energy Prices," working paper, The University of Kansas.

Popp, David (1997), *Induced Innovation, Energy Prices, and the Environment*, Ph.D. Thesis, Yale University.

Putnam, Jonathan (1996), *The Value of International Patent Protection*, Ph.D. Thesis, Yale University.

Rosenberg, Nathan (1994), *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, Fredric M. (1982), "Interindustry technology flows and productivity growth," *Review of Economics and Statistics*, 627-34.

Scherer, Frederic M. (1984), "Using Linked Patent and R&D Data to Measure Interindustry Technology Flows," in Zvi Griliches, ed., *R&D, Patents and Productivity*, University of Chicago Press, Chicago, 417-461.

Thirtle, Colin G. and Vernon W. Ruttan (1987), *The Role of Demand and Supply in the Generation and Diffusion of Technical Change*, Harwood Academic Publishers, London.

Trajtenberg, Manuel, 1990, "A Penny for Your Quotes: Patent Citations and the Value of Innovations," *Rand Journal of Economics*, 21, 172-187.

U.S Department of Commerce (1994), Patents and Trademark Office, *Index of Patents, Part II: Index to Subjects of Inventions*, U.S. Government Printing Office, Washington, D.C.

ENDNOTES

[1] Examples of the first type of study include Griliches and Mairesse (1984), Clark and Griliches (1984) and Scherer (1982, 1984). See Griliches (1995) for a survey of this work. Examples of the second type of work include Putnam (1996), Lanjouw (1993), Pakes (1986), and Pakes and Schankerman (1984). See Lanjouw, Pakes, and Putnam (1996) for a review of this literature.

[2] See, for example, Binswanger (1978) and Thirtle and Ruttan (1987).

[3] It is important to note that not all inventions are patented. Patents provide monopoly protection for a period of years (17 in the United States) in return for making the invention public knowledge. If an inventor feels that making an invention public will enable others to "invent around" the patent in such a way as to improve the existing invention but not violate the patent, the inventor may prefer to keep the invention secret.

[4] Empirical work demonstrating the effect of prices or regulation on innovation include Popp (1998), Newell, Jaffe and Stavins (1998), and Jaffe and Palmer (1997). Theoretical models include Milliman and Prince (1989, 1991) and Krutilla and Jung (1996).

[5] "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.

[6] In fact, patent citations are more useful to researchers than journal citations. Because of the detailed patent examination process, and because of the legal importance placed on patent citations, we can be more confident that patent citations are an accurate representation of the knowledge upon which an inventor built, rather than just, for example, a random survey of previous literature in a journal article.

[7] A more detailed presentation of the search model is provided in appendix A of this paper.

[8] This is also the intuition behind the evolutionary theory of invention, discussed in Nelson and Winter (1982).

[9] An analogy with academic research might be useful here. One of the papers cited most often in economic literature is Ronald Coase's 1960 treatise on social cost. This paper is cited often because it motivated several papers to examine his ideas more closely. Thus, Coase's article shifted the distribution of research related to social cost outward. Now, consider what would happen if an author came along and wrote a paper that proved beyond doubt that Coase was correct (or that he was incorrect). This paper would be of great value to economists, so that it would be far to the right of the distribution of possible outcomes. However, if the paper truly settled the debate once and for all, there would be little work to follow it, so that the paper would not be cited often. The distribution of outcomes would not shift. This would be an example of invention potential exhaustion.

[10] In the search model of R&D, the only result that matters is that of the most successful search. Because of this, the theory of order statistics demonstrates that the distribution of outcomes shifts outward as more searches are taken. Thus, the distributions in figure 3 should thus be thought of as distributions for a certain number of searches. If the number of searches changes, the nature of the distributions will also change.

[11] Since institutional changes will affect all patents equally, this parameter is not indexed by *i*.

[12] 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 (1995) as well. 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.

[13] 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. Thus, α_1 represents citation practices in 1974-1975, α_2 citation practices in 1976-1977, etc.

[14] To verify that the upswing is not simply due to truncation of the data set, the regressions were rerun with two years less data (that is, cited patents through 1987 and citing patents through 1989). The upswing is not present in the results of this regression.

[15] As this discussion suggests, the ideal specification of the model would use the number of patents, rather than time, to estimate the decay rate. Unfortunately, the model does not converge when specified in this manner.

[16] Regressions were also run for renewals in the fourth and twelfth years, and similar results were obtained. However, the results of the regression for renewals after eight years are best for two reasons: 1) after 4 years, patents that are highly cited have not received many of their citations, and 2) there are few patents in the data set that are old enough to have needed renewal after twelve years.

[17] See, for example, Griliches (1989, 1990) and Evenson (1991).

[18] More correctly, the issue in question is whether the impact of invention potential exhaustion is greater than the impact of recharge.

[19] Note that, if individual regressions are run for each technology group, the coefficient on the number of patents is only positive for 5 of the 12 technologies, so the sign should be interpreted carefully. However, when it is negative, it is rarely significant.

[20] For the solar energy ratio, the sum of patents in all four solar energy patents is used. For the demand-side technologies, patents for waste heat, heat exchange, heat pumps, Stirling engines, and continuous casting are used.

[21] For example, refrigeration is energy-intensive activity that could be improved by R&D. Such R&D might focus on more efficient compressors or better insulation materials. However, patents in the refrigeration subclass are organized by the end-use of refrigeration, rather than by the components of refrigerators. Since there are other reasons for inventive activity in refrigerators, such as to reduce CFC emissions, refrigeration patents are not used.

Figure 1 -- Ratio of Patents to R&D Spending

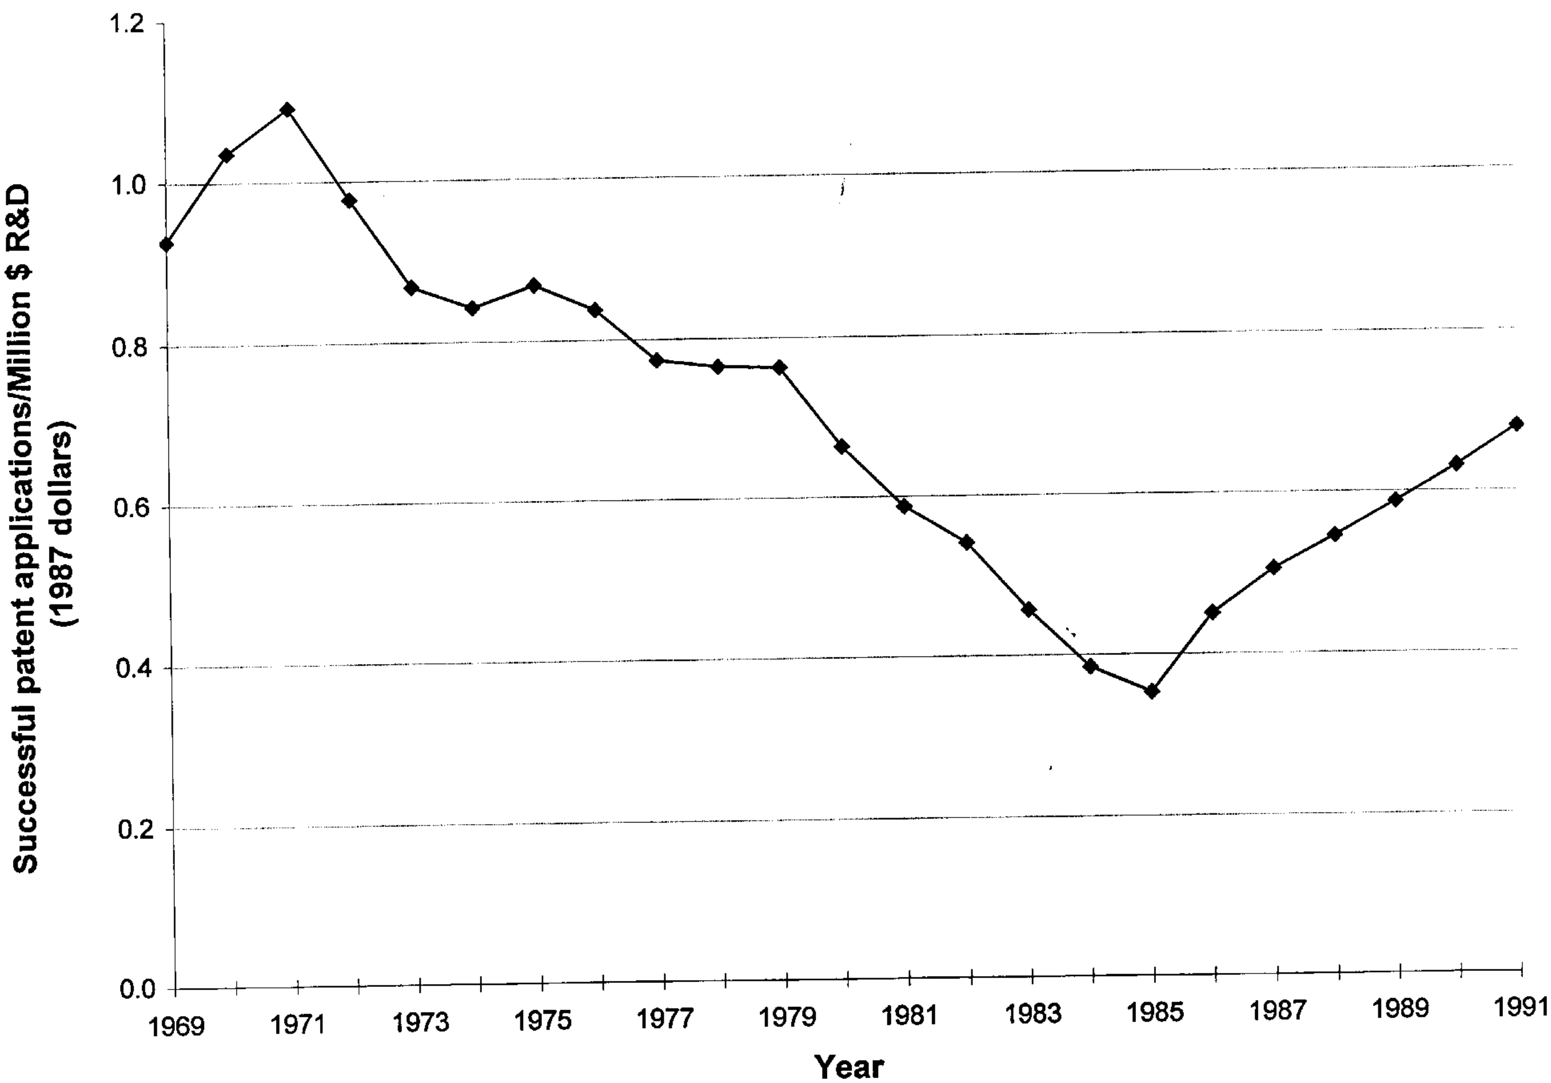

Figure 2 -- The Expected Value of Research

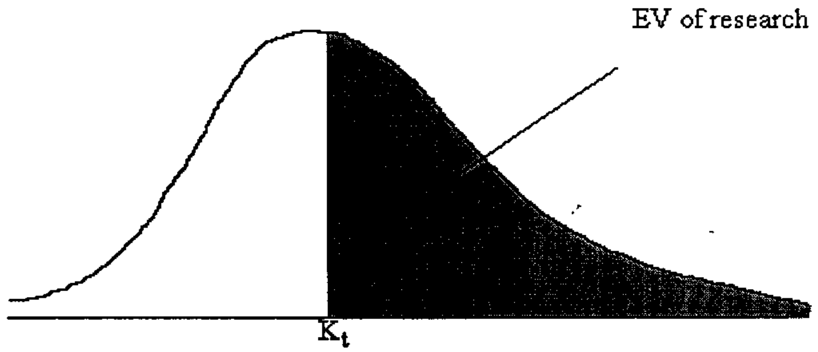

The expected value of research is the area under the distribution of potential outcomes to the right of the current level of knowledge, K_t.

Figure 3 – Invention Potential Exhaustion and Recharge

A.

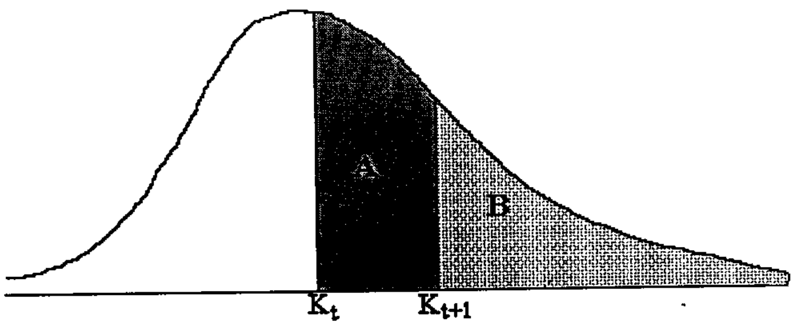

Invention Potential Exhaustion: The expected value of research in period t is areas A and B, but only area B in period t+1.

B.

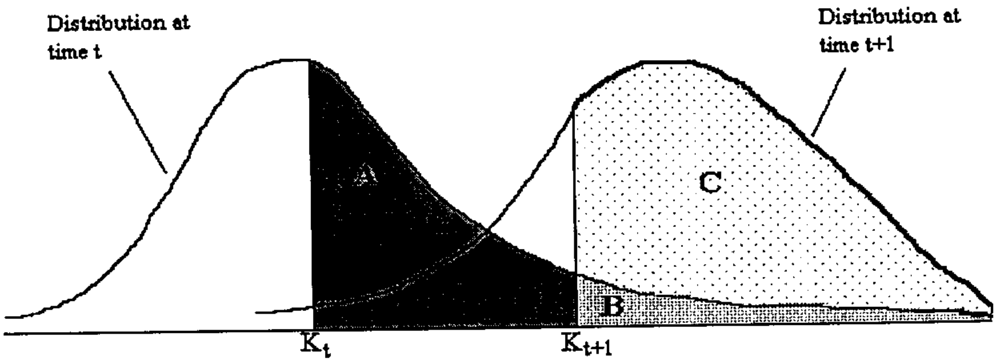

Recharge: The expected value of research at time t is A + B. At time t + 1, it is B + C. Since area C > area A, recharge increases the expected value of research in period t + 1.

Figure 4 -- Citations and the Expected Value of Research

Project 1:

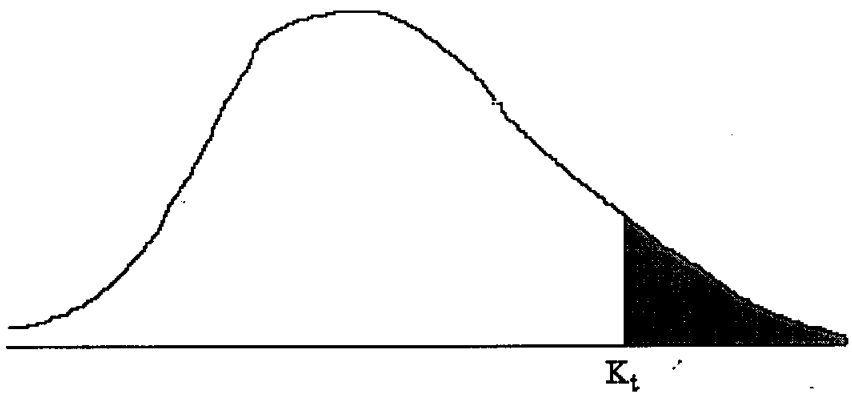

Project 2:

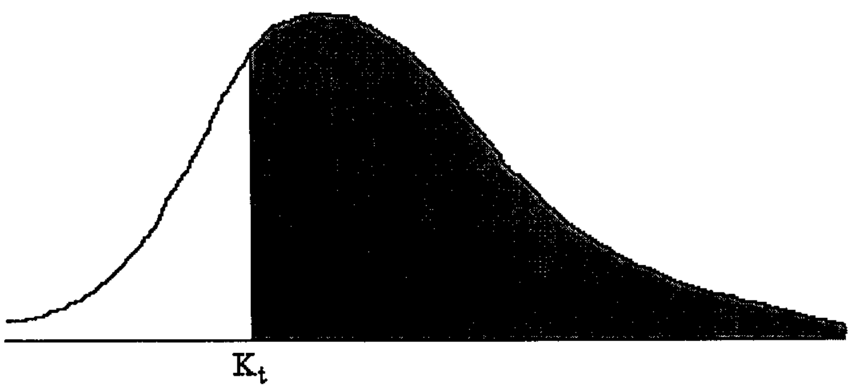

The expected value of research is greater for project 2 than project 1. Mre new inventions will follow from projectd 2, resulting in more citations to the patent representing K_t.

Figure 5 -- Probability of Citation

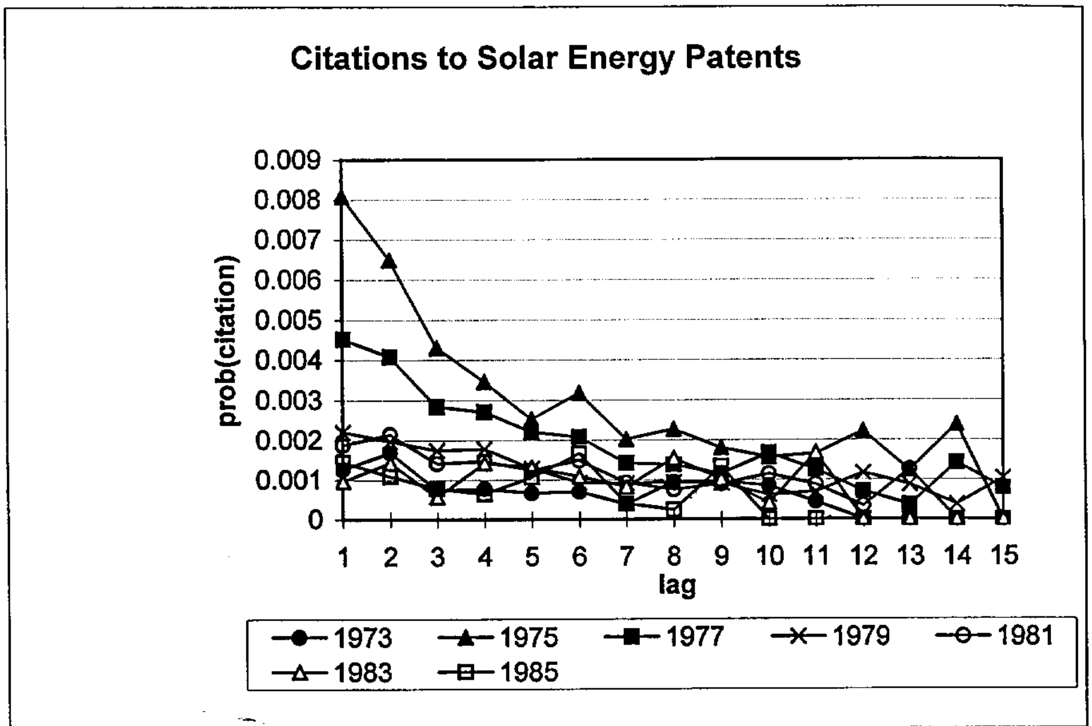

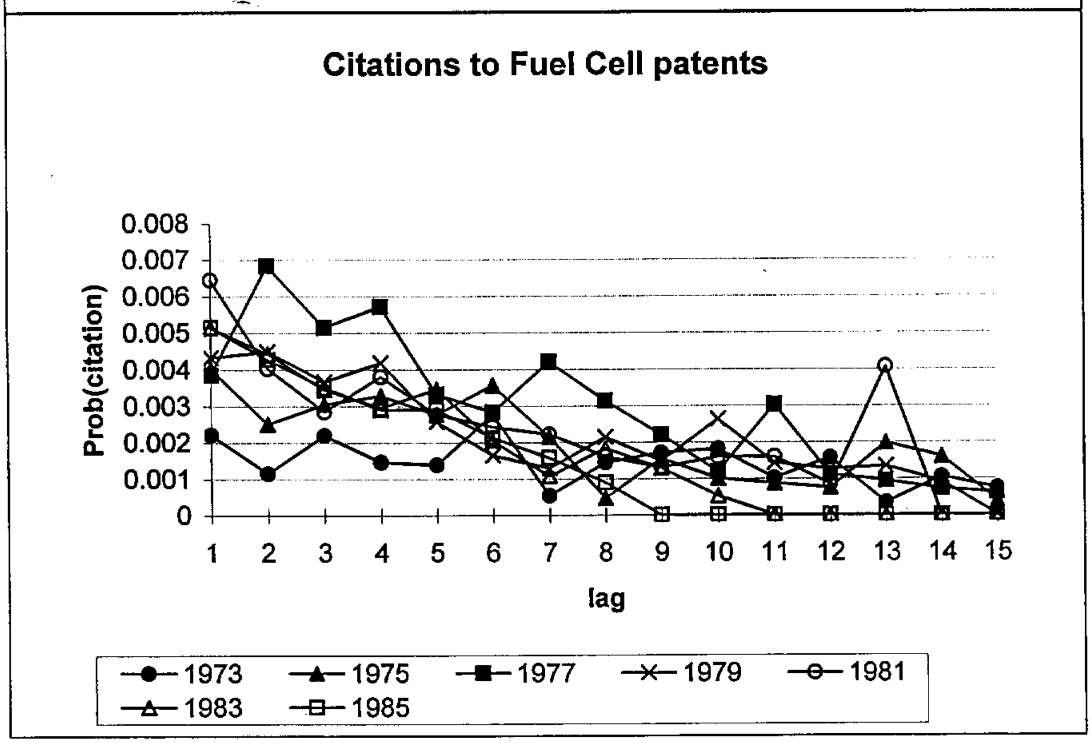

The figure presents the probability that patents granted in year x will be cited by patents applied for in following years. Each line represents the patents granted in a different year. The x-axis is the number of years since the patent was granted.

Figure 6 -- Productivity Estimates: Citations to Patents in Same Technology Group

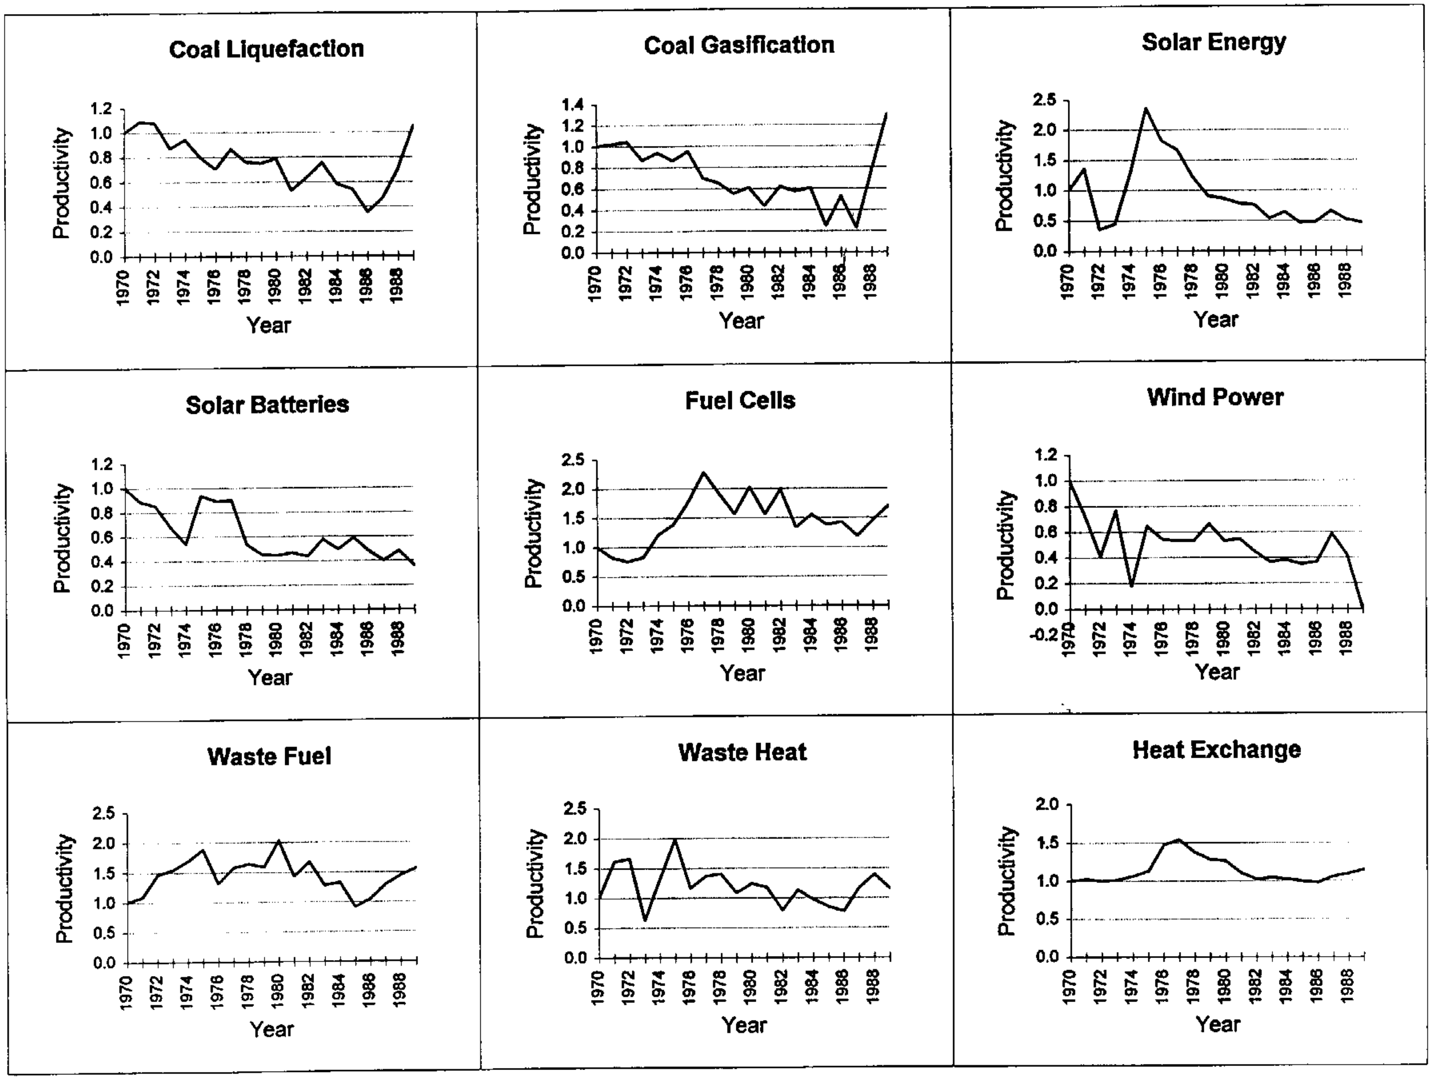

Figures show the productivity estimates for each technology group, with 1970 normalized to 1 in each case. Note that, for most technologies, there is a declining trend to the estimates, suggesting diminishing returns to research over time. This will be explored more carefully in section V.

Figure is continued on the next page.

Figure 6 -- Productivity Estimates: Citations to Patents in Same Technology Group

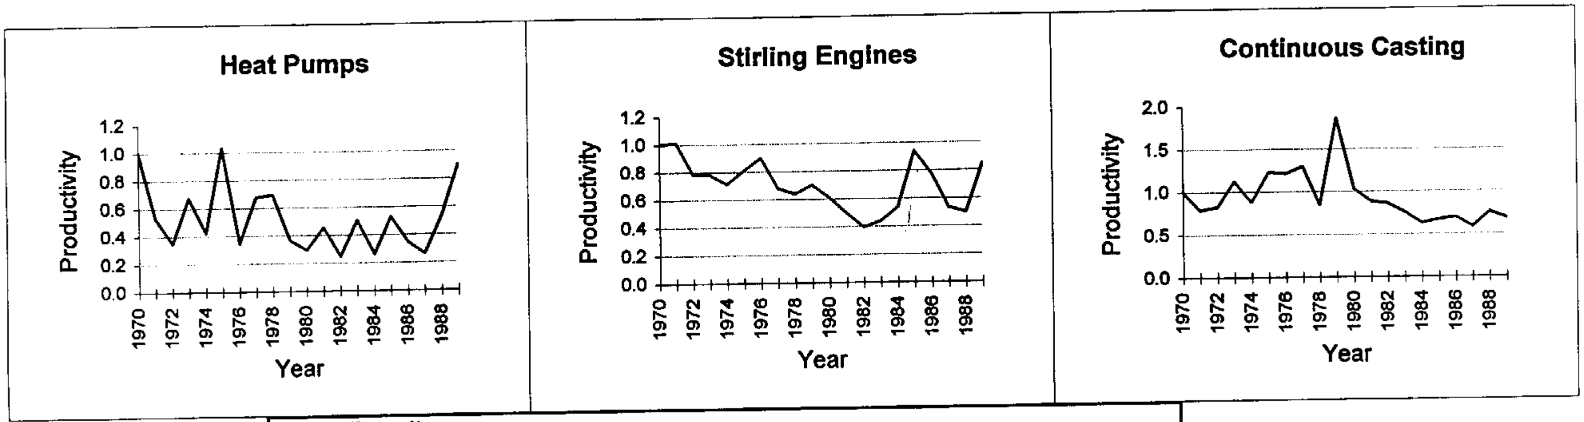

Other Results

citing practices:	estimate:	standard error:	**group dummies:**	estimate:	standard error:
1974-1975	1.000	N/A	coal liqefaction:	4.621	1.144
1976-1977	0.906	0.056	coal gasifcation:	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
			heat pumps:	10.362	3.443
decay:	0.392	0.018	Stirling engines:	5.926	1.772
diffusion:	0.00251	0.00061	continuous casting:	1.000	N/A
adjusted R-square:		0.778			

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. This will be explored more carefully in section V.

Figure 7 -- Productivity Estimates: Citations to All Patents

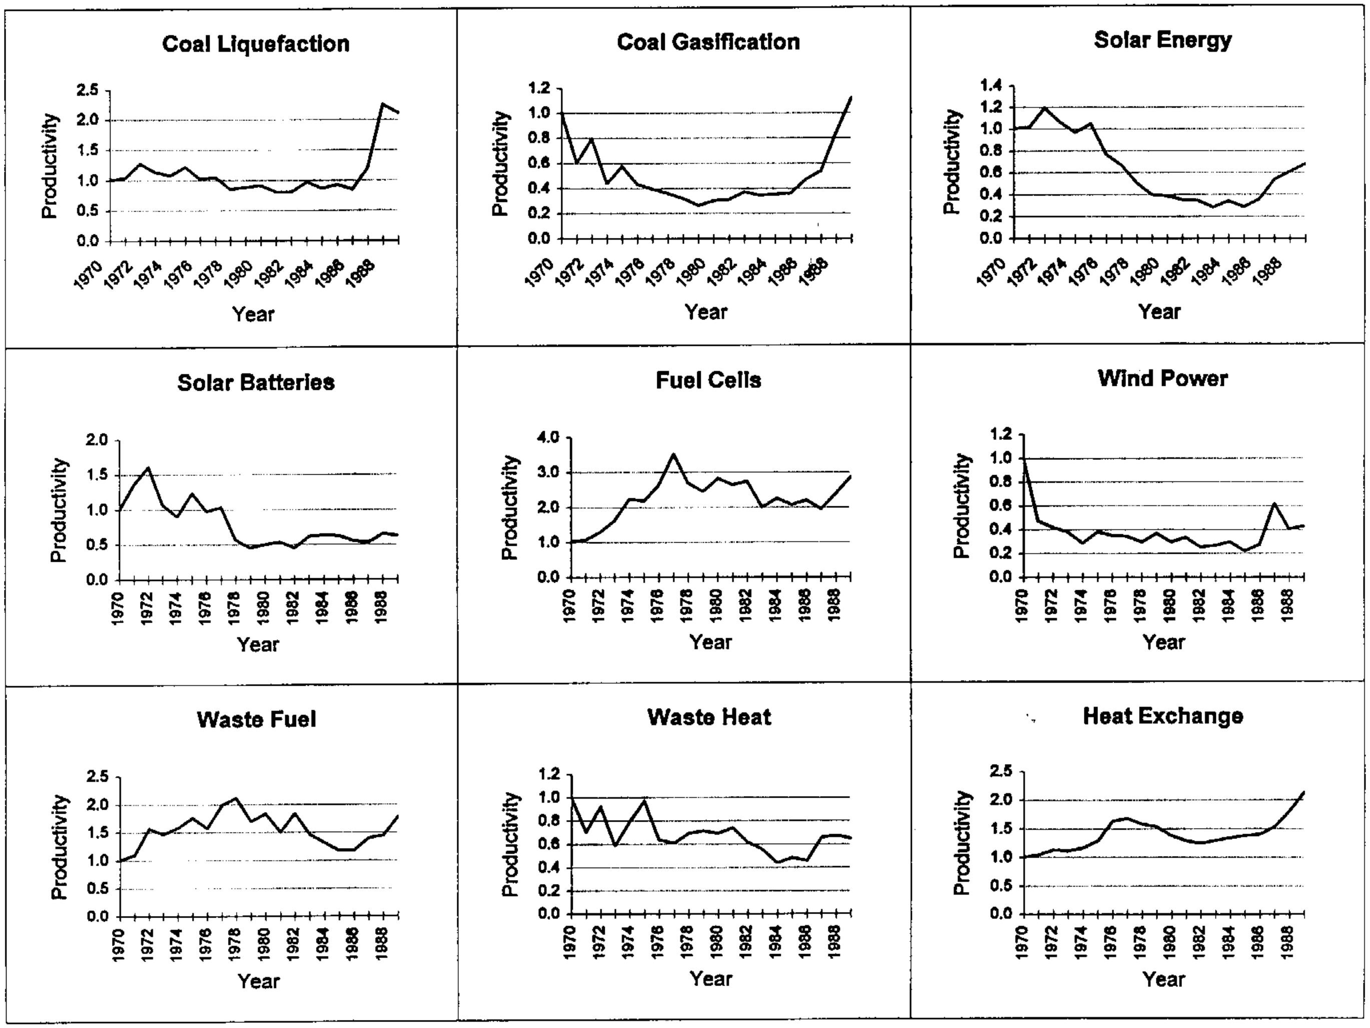

Figures show the productivity estimates for each technology group, with 1970 normalized to 1 in each case. Although there is still some evidence of diminishing returns, the decline in productivity estimates is not as pronounced as in the regression that used only citations in the technology group. Also note that many groups exhibit a U-shaped pattern. Explanations for this will be explored in section V.

Figure is continued on the next page.

Figure 7 -- Productivity Estimates: Citations to All Patents

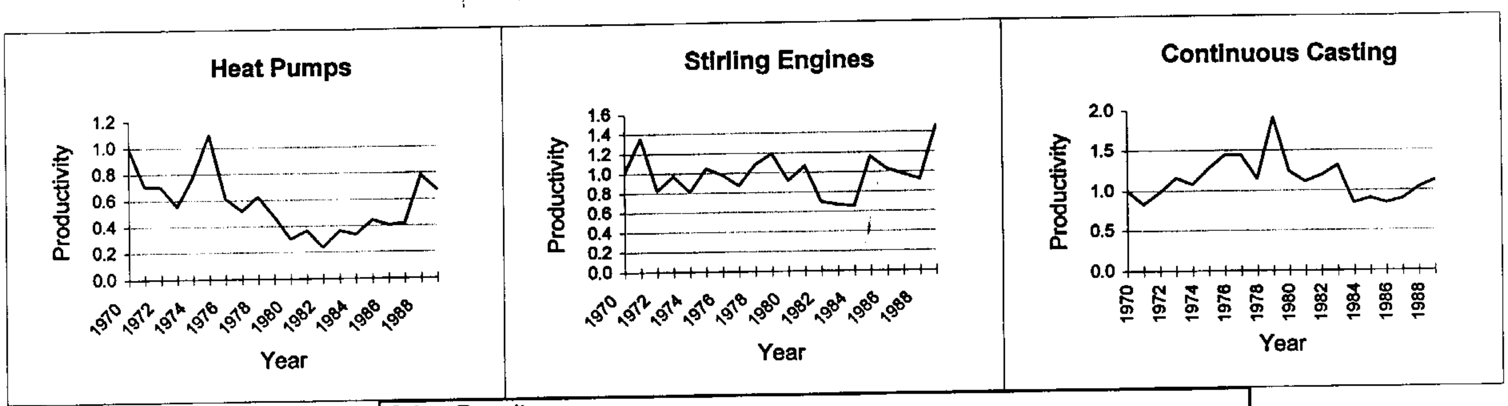

Other Results

citing practices:	estimate:	standard error:	group dummies:	estimate:	standard error:
1974-1975	1.000	N/A	coal liqefaction:	4.654	1.212
1976-1977	0.878	0.047	coal gasifcation:	14.400	3.773
1978-1979	0.807	0.062	solar energy:	2.914	0.872
1980-1981	0.744	0.079	solar batteries:	3.164	0.894
1982-1983	0.764	0.103	fuel cells:	1.219	0.419
1984-1985	0.789	0.132	wind:	27.312	7.596
1986-1987	0.771	0.153	waste fuel:	2.445	0.790
1988-1989	0.790	0.182	waste heat:	4.786	1.645
1990-1991	0.733	0.192	heat exchange:	0.225	0.076
			heat pumps:	21.497	5.766
decay:	0.320	0.016	Stirling engines:	6.087	1.905
diffusion:	0.00245	0.00062	continuous casting:	1.000	N/A
adjusted R-square:		0.800			

Figures show the productivity estimates for each technology group, with 1970 normalized to 1 in each case. Although there is still some evidence of diminishing returns, the decline in productivity estimates is not as pronounced as in the regression that used only citations in the technology group. Also note that many groups exhibit a U-shaped pattern. Explanations for this will be explored in section V.

Figure 8 -- Productivity Estimates and Patent/R&D Ratios

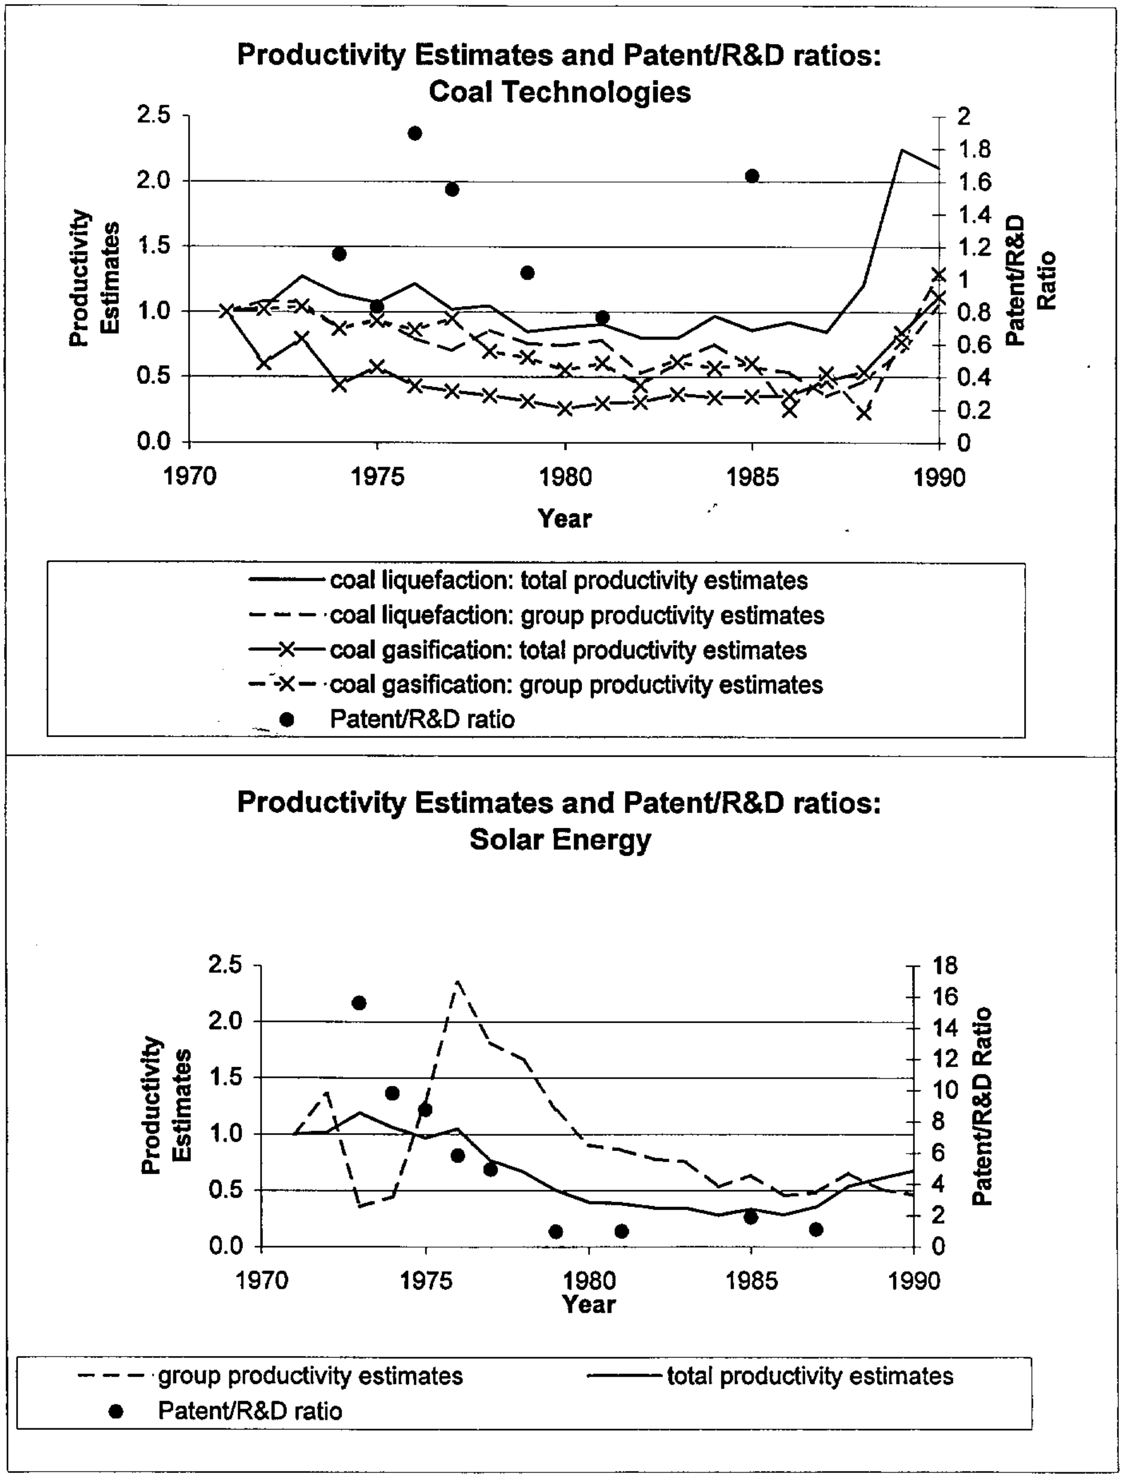

Lines in the figures show values of the productivity estimates, lagged one year, for the technologies listed. The dots show patent-to-R&D ratios. Note that lagged values of the productivity estimates provide a good explanation of changes in the patent-to-R&D ratio.

Figure continued on the next page.

Figure 8 -- Productivity Estimates and Patent/R&D Ratios

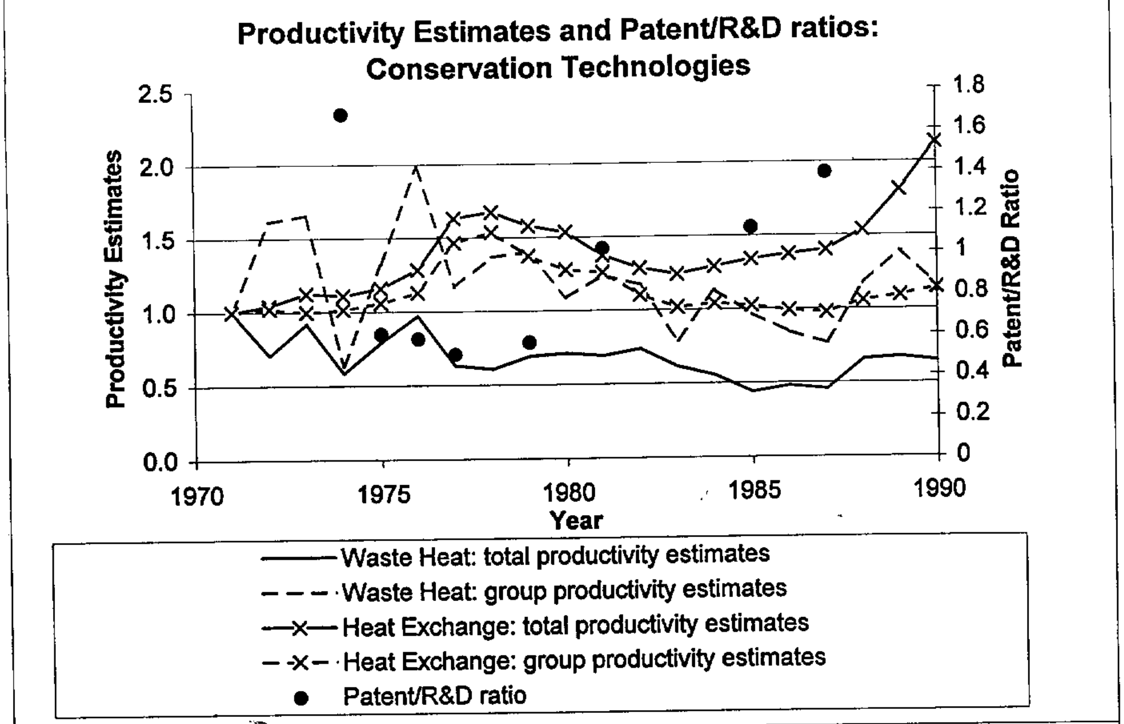

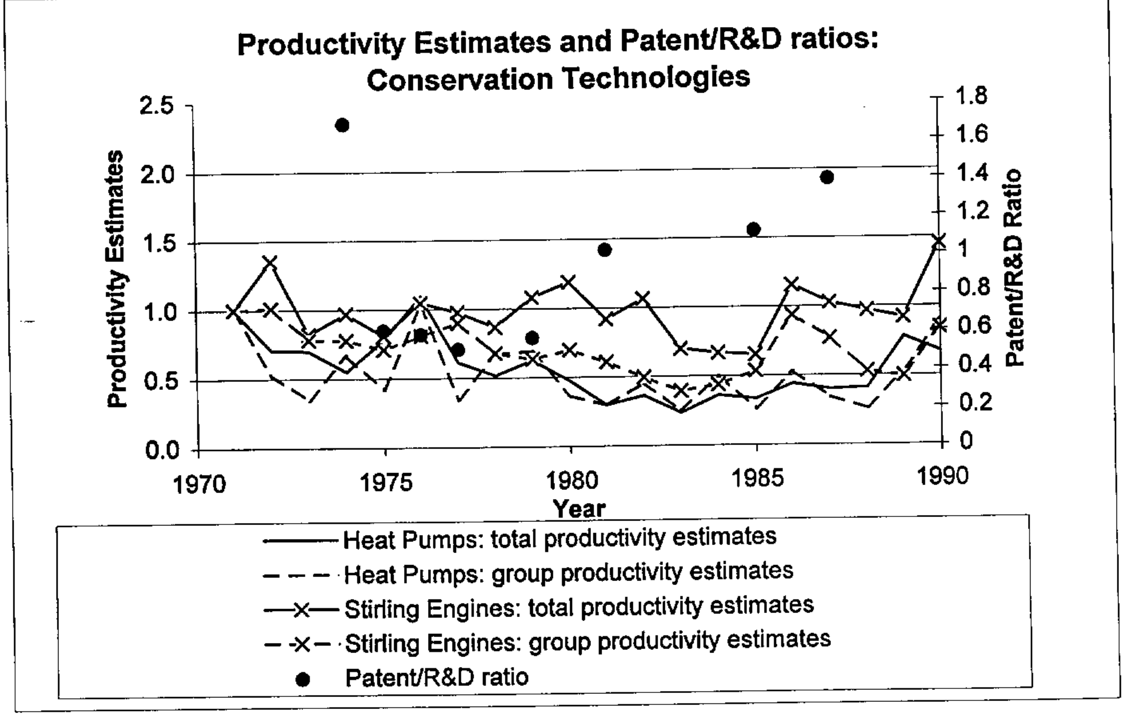

Lines in the figures show values of the productivity estimates, lagged one year, for the technologies listed. The dots show patent-to-R&D ratios. Note that lagged values of the productivity estimates provide a good explanation of changes in the patent-to-R&D ratio.

Figure A1 -- How Research Affects the Distribution of Possible Outcomes

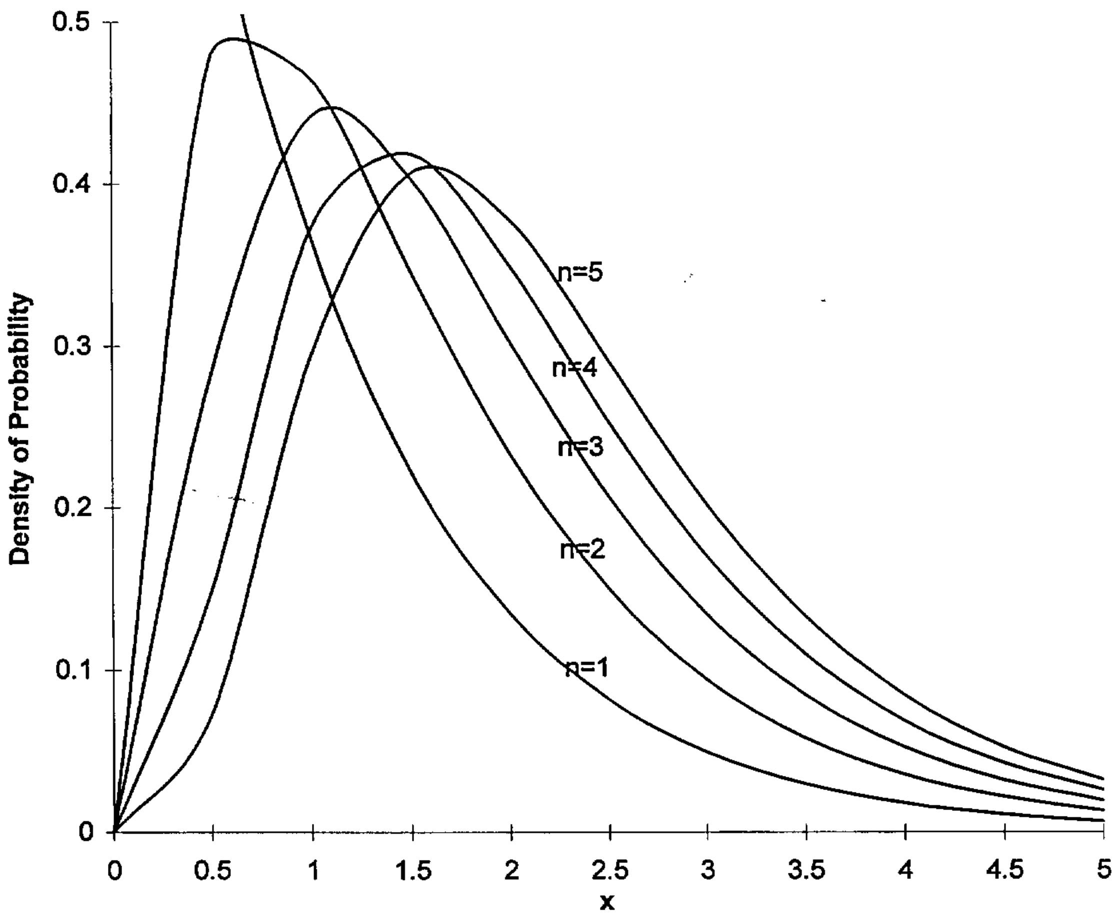

The figure shows how the distribution of potential outcomes shifts outward as the number of trials increases. The function illustrated is the exponential variate. It is:

$$j_n(z) = n(1-e^{-x})^{n-1}e^{-x}$$

Source: Evenson & Kislev (1975)

Figure B1- Patent Applications by U.S. Inventors

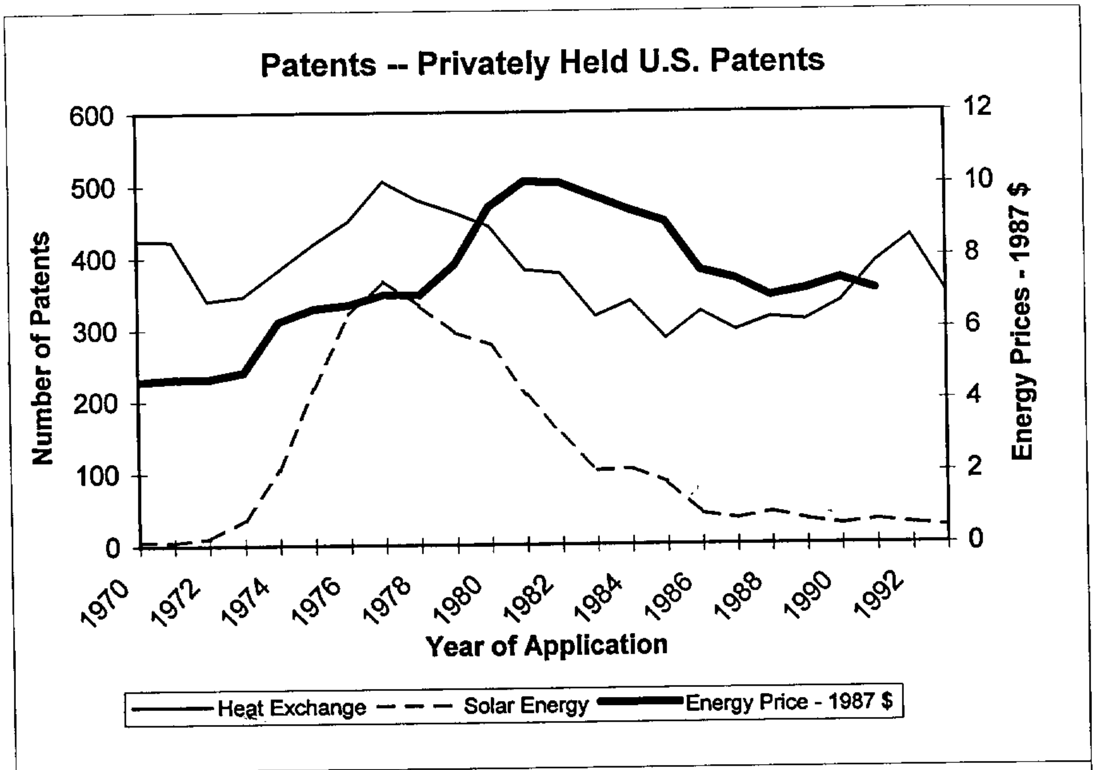

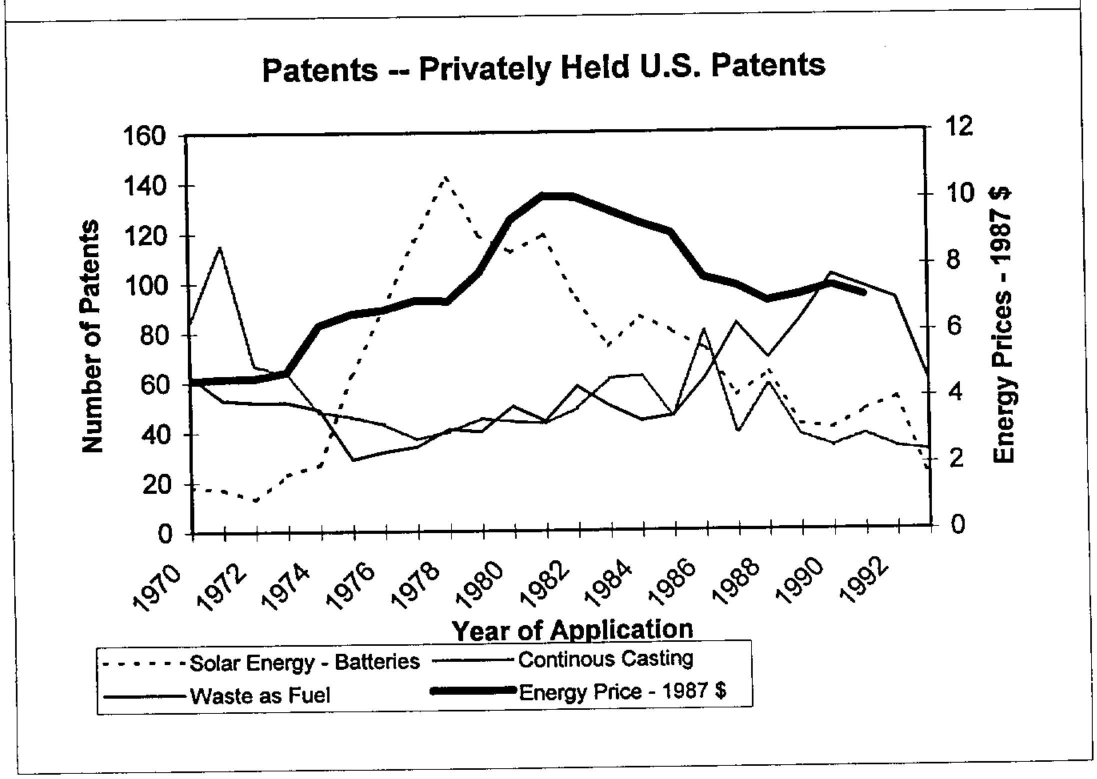

NOTES: Energy prices are the cost per million Btu of energy consumption, in 1987 dollars.

Figure is continued on the next page.

Figure B1 -- Patent Applications by U.S. Inventors (continued)

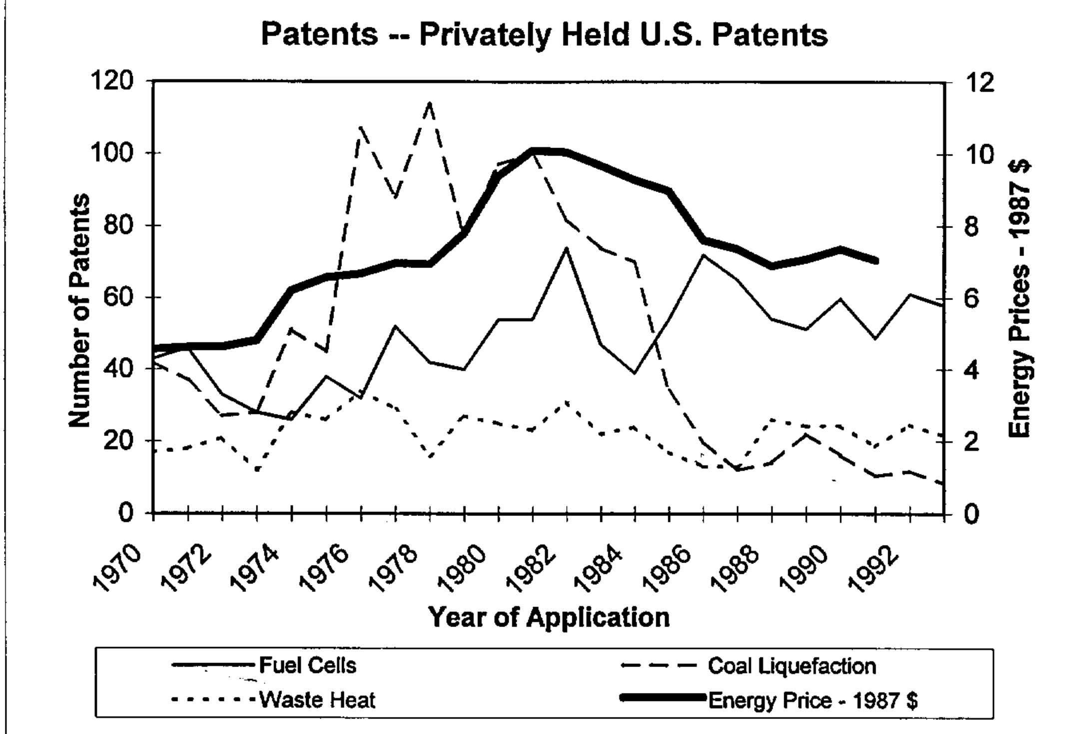

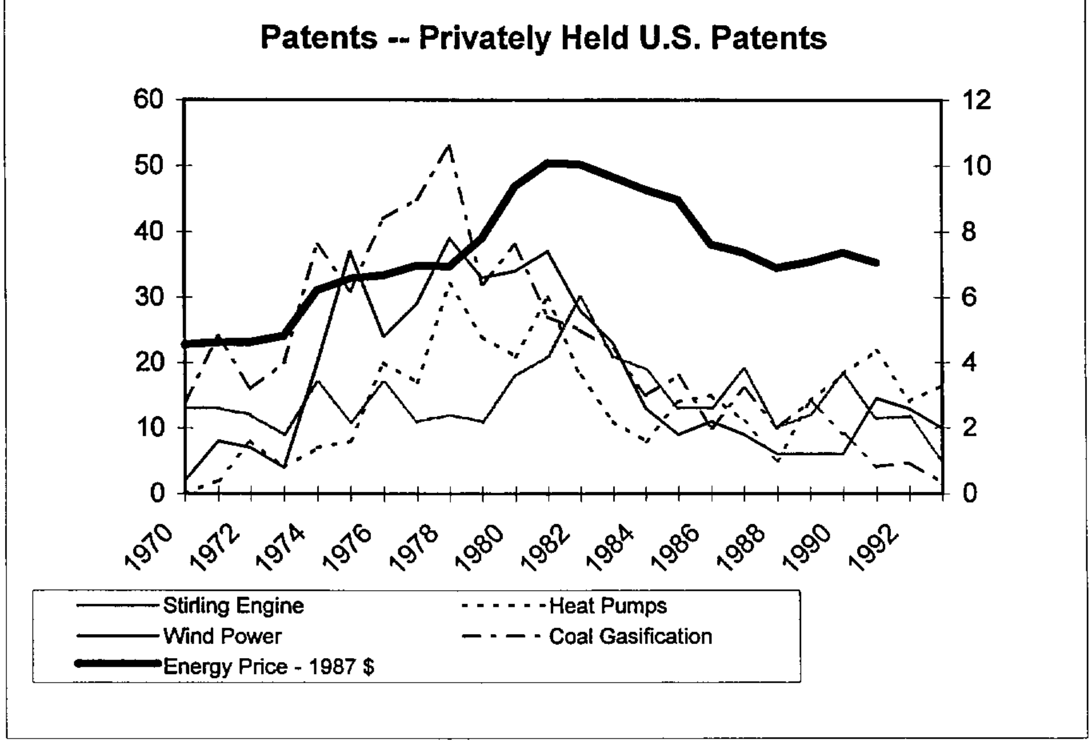

NOTES: Energy prices are the cost per million Btu of energy consumption, in 1987 dollars.

Table 1 -- Marginal Effects of an Additional Citation on the Probability of Renewal

Data after eight years of a patent's life

Technology Group	*Δ probability of renewal*
Coal Gasification	0.059 **
Coal Liquefaction	0.042 ***
Continuous Casting	0.037 ***
Fuel Cells	0.015 ***
Heat Pumps	0.125 ***
Heat Exchange	0.031 ***
Solar Batteries	0.033 ***
Solar Energy	0.019 *
Stirling Energy	0.045 **
Wind Power	0.023
Waste as Fuel	0.012 **
Waste Heat	0.046 ***
Mean	0.041
Median	0.035

* - estimated coefficient is significant at the 10% level
** - estimated coefficient is significant at the 5% level
*** - estimated coefficient is significant at the 1% level

Table 2 -- Test of Diminishing Returns in the Productivity Parameters

Dependent variable: Productivity estimates

	(eq. 5)		(eq. 6)	
Independent Variables	Within Group	Total Citations	Within Group	Total Citations
Intercept	1.071	0.923	1.028	0.966
	(14.429)	(14.115)	(14.246)	(13.282)
Time	-0.011	0.004	-0.011	0.004
	(-1.821)	(0.769)	(-1.913)	(0.648)
Number of patents granted in the cited year	--	--	0.0002	-0.0006
	--	--	(2.910)	(-5.418)
Adjusted R-square	0.014	0.003	0.046	0.111

Table 2 shows the results of tests for diminishing returns in the productivity estimates. Equation 5 regresses the productivity estimates on a time trend, while equation 6 regresses the productivity estimates on both time and the number of patents granted in the cited year. The data from each technology group was pooled to obtain these estimates. The significant negative coefficient on the time trend is evidence of diminishing returns when only using within group citations. However, when citations to all patents are used, diminishing returns occur within a given year, rather than across time. Reasons for the differing results are discussed in the text.

Table 3 -- Patents/R&D Expenditures

R&D Expenditures in millions of 1987 dollars

Year	Solar R&D	Solar Energy Patents/R&D	Other Coal R&D	Coal Patents/R&D	Conservation R&D	Conservation Patents/R&D*
1973	5	15.61	N/A	N/A	N/A	N/A
1974	15	9.81	84	1.15	295	1.69
1975	37	8.78	99	0.83	858	0.61
1976	80	5.86	80	1.89	985	0.59
1977	111	4.98	89	1.55	1,197	0.51
1978	N/A	N/A	N/A	N/A	N/A	N/A
1979	482	0.98	114	1.04	1,030	0.57
1980	N/A	N/A	N/A	N/A	N/A	N/A
1981	372	1.00	177	0.77	505	1.02
1982	N/A	N/A	N/A	N/A	N/A	N/A
1983	N/A	N/A	N/A	N/A	N/A	N/A
1984	N/A	N/A	N/A	N/A	N/A	N/A
1985	102	1.90	33	1.64	342	1.12
1986	N/A	N/A	N/A	N/A	N/A	N/A
1987	95	1.12	N/A	N/A	281	1.39

N/A: R&D expenditures not available for this year

R&D data taken from National Science Foundation Survey of Industrial Research and Development

** - Patents included in conservation category: Waste Heat, Heat Exchange, Heat Pumps, Stirling Engines, & Continuous Casting*