# Identifying the Age Profile of Patent Citations: New Estimates of Knowledge Diffusion *

Aditi Mehta
*U.S. Department of Justice*

Marc Rysman
*Department of Economics, Boston University*

Tim Simcoe
*J.L. Rotman School of Management, University of Toronto*

September 29, 2008

*The views expressed are not purported to reflect those of the U.S. Department of Justice. We thank Ajay Agrawal, Bronwyn Hall, Josh Lerner, David Mowery, Scott Stern, Kevin Lang and several seminar audience for helpful comments. The editor and three anonymous referees provided helpful advice in improving the paper. Address for correspondence: Joseph L. Rotman School of Management, 105 St. George Street, Toronto, ON M5S 3E6,Canada. E-mail: timothy.simcoe@rotman.utoronto.ca

**Abstract**

Previous research studies the age-profile of patent citations to learn about knowledge flows over time. However, identification is problematic because of the collinearity between application-year, citation-year, and patent-age. We show empirically that a patent's "citation clock" does not start until it issues, and propose a highly flexible identification strategy that uses the lag between application and grant as a source of exogenous variation. We examine the potential bias if our assumptions are incorrect, and discuss extensions into other research areas. Finally, we use our method to re-examine prior results on citation age-profiles of patents from different technological fields and application-year cohorts. JEL Codes: L1, O3

Keywords: age profile, grant lag, knowledge diffusion

# 1 Introduction

Patent citations are widely used in empirical work as an observable proxy for knowledge flows. Since the pioneering research of Jaffe, Trajtenberg & Henderson (1993), a large number of papers have used citations to study the geographic localization of knowledge spillovers. Another line of research uses citations to measure the economic or technological significance of a given patent.[1] In both cases, researchers are often interested in the distribution of citations received over the life of a patent. The shape of this citation "age profile" may reveal information about the rate of knowledge utilization or the duration of technology life-cycles.

Unfortunately, empirical research on the age-profile of patent citations quickly encounters a well-known identification problem. Suppose that a patent is born in year $b$, and consider the number of citations it receives in calendar year $t$ conditional on its age $a$. If the birth-year, citing-year, and age all influence the citation rate $C_t$, our model should include each of these effects: $C_t = f(b, t, a)$. However, since the patent's age in year $t$ is usually defined as $a = t - b$, we have $C_t = f(b, t, t - b) = \hat{f}(b, t)$. In other words, it is impossible to identify all of the parameters in this model given the deterministic relationship between age, birth-year and citing-year.

Prior research has established the importance of controlling for both birth-year and citing-year effects when estimating a citation-age distribution (Hall, Jaffe & Trajtenberg, 2002). Controlling for the birth-year of the cited patent (called the "cohort effect") is important because the standards for granting a patent have changed over time due to policy changes and funding issues at the United States Patent and Trademark Office (USPTO). [2] Controlling for the birth-year of the citing patent (called the "calendar effect") is important because patents have received more citations in recent years, reflecting an overall increase in patenting (Kortum & Lerner, 1999) and the adoption of computerized search at the USPTO. Given these constraints, prior research has relied on functional form assumptions to identify the age-distribution. A particularly common approach[3] is to specify a double-exponential age-profile: $C_t = f(b, t)e^{\beta_1 a}(1 - e^{\beta_2 a})$.

[1] Papers concerned with knowledge spillovers include Mowery, Oxley & Silverman (1996); Jaffe & Trajtenberg (1999); Agrawal, Cockburn & McHale (2006); Thompson & Fox-Kean (2005); and Alcacer & Gittelman (2006). Papers focused on patent value include Trajtenberg (1990); Albert, Avery, Narin & Mcallister (1991); Harhoff, Narin, Scherer & Vopel (1999); or Hall, Jaffe & Trajtenberg (2005).

[2] Controlling for the cohort effect also controls for changes in the quality of innovations over time. As discussed in Hall et al. (2002), one may or may not want to control for this issue depending on the application.

[3] Applications of this method include Jaffe & Trajtenberg (1996), Hall et al. (2002); Bacchiocchi & Montobbio (2004); Branstetter & Ogura (2005); Adams, Clemmons & Stephan (2006); and Marco (2007).

In this paper, we introduce a new method that can identify all three effects – age, calendar and cohort – using a highly flexible citation-generating function. We use this method to revisit results from Hall et al. (2002) on variation in the rate of knowledge diffusion across industries. We also address some new questions, such as whether diffusion begins at application or grant (see also Johnson & Popp, 2003), and how diffusion rates have changed over time.

The central idea is novel, yet extremely simple. We exploit the patent review process as a source of exogenous variation in the "age" of patents within a given cohort. That is, we assume that there is a random lag between the application year and grant year of a patent and that the age process begins when a patent is granted rather than at application time. Formally, we continue to define the birth year $b$ as the application year but we redefine age as $a = t - g$ (rather than $a = t - b$), where $g$ is the grant year. Since we observe patents from the same birth cohort with different "ages" in any given calendar year, $C_t = f(b, t, a)$ is identified.

Our argument for this identification strategy is based on three observations. First, the appropriate definition of a patent's birth cohort is the application year. While some papers have used the grant year, Hall et al. (2002) definitively favor the application year as capturing the relevant economic information:

> The actual timing of patented inventions is closer to the application date than to the (subsequent) grant date... Indeed, the mode of operation at the Patent Office underwent significant changes in the past decades, thereby introducing a great deal of randomness (which has nothing to do with the actual timing of the inventions) into any patent time series dates by grant year. Thus, and whenever possible, the application date should be used as the relevant time placer for patents. (pg. 409-410)

Our second observation is that a patent's "citation clock" does not appear to start until after it has been granted. This makes sense because the majority of U.S. patents are published and become publicly available on the grant date. To be clear, by assuming that the grant date marks the beginning of the age process, we ignore the potential importance of pre-grant diffusion in determining the age-profile of citations. However, we present statistical evidence that strongly suggests the importance of the grant date in determining the timing of citations. In particular, when visually comparing the age-profile of patents from the same application year, it appears that those with a longer grant lag have an identical age profile shifted over by one year. This finding is relevant whenever a researcher must choose an appropriate "birth date" (regardless of whether they are interested in the citation age profile).

Our third observation is that several factors lead to a great deal of random noise in the grant or pendency lag; the time period between filing an application and receiving a patent. For example, Cockburn, Kortum & Stern (2002) suggest that the identity of the patent examiner influences how quickly an application is processed within the USPTO. Processing speed may also be influenced by random events, such as employee turnover or short-run fluctuations in the volume of applications. We use the resulting variation to "solve" the age-year-cohort identification problem, and obtain new insights about the shape of the age profile in the absence of functional form assumptions.

Of course, the proposed identification strategy rests on the assumption that the grant lag is exogenous to citations. This is a potential concern, since Johnson & Popp (2003) and Popp, Juhl & Johnson (2004) argue that longer pendency lags are associated with increased citations.[4] We cannot control for the grant lag directly because it re-introduces the original collinearity problem (i.e. $t - b + l = a$, where $l$ is the grant lag). However, if we assume the grant lag has a positive impact on citations, then our procedure will estimate an upper bound on the true age coefficients. Intuitively, the omitted grant lag is negatively correlated with age by construction, and potentially positively correlated with citations, leading to coefficients that are biased towards zero. We test the importance of excluding the grant lag and show that our main estimates are likely to be close to the true parameters. We also describe an instrumental variables strategy that would address grant lag endogeneity and perhaps extend our method to a broader range of applications, such as the study of individual productivity over the life-cycle, as in Hall, Mairesse & Turner (2007) or Jones (2007).

A separate benefit of our approach is that we can flexibly estimate the age profile at age zero. The functional forms utilized in previous work such as Hall et al. (2002) typically cannot handle age zero and so those authors would simply assume that there were no citations in the birth year. We show below that this assumption is restrictive. Our approach can also estimate the age profile for negative ages, which are a natural occurrence in our approach as we discuss below.

After describing and testing our identification strategy, we use it to revisit several published results on the age-profile of patent citations. We begin by estimating a complete set of application-year, citing-year, and age-since-grant effects for all patents in the NBER U.S. patentdata set, and comparing these estimates to an age-profile obtained by fitting a double exponential function. While both profiles have a similar strongly peaked shape, our approach

[4] Formally, they study the length of time between the priority date and the grant date, which is greater than what we term the grant lag (see the discussion in the next section), but many issues are likely to carry over.

finds that peak citation rates occur somewhat earlier, one year after grant rather than four years after application. We then use our technique to revisit Hall et al. (2002)'s study of the difference in citation profiles across industry sectors. While a comparison of mean citation ages across industries reveals a slightly different ordering (using our method, computers and electronics receive the earliest citations and pharmaceuticals the latest), the overall results are remarkably similar. We conclude that the double-exponential specification provides a good approximation to the non-parametric age distribution.

Finally, we take an initial look at how the citation age profile has changed over time by allowing the age effects in our model to vary across application-year cohorts. Despite important changes in the technology for diffusing ideas, the legal treatment of patents, and arguably the culture of patenting and patent enforcement, we find that the citation age profile has remained remarkably stable. While there is some evidence that the most recent patents have experienced a slower depreciation rate, more data would be preferable to draw firm conclusions.

This paper offers a new approach to a particular version of the age-year-cohort identification problem. Hall, Mairesse & Turner (2007) offer a more general discussion of this issue, which appears in a variety of fields (e.g when estimating life-cycle productivity or hedonic pricing models). For instance, Galenson & Weinberg (2000) study the age profile of output among modern American painters. Hall et al. (2007) cite a number of papers that rely on functional form assumptions to overcome the identification problem and propose a general strategy for testing the statistical significance of any given set of effects. They also point out that fixed-effects estimation typically aggravates the problem—an issue that we do not address here.

Given the large amount of empirical research that relies on citation data we believe our method could be used in a wide variety of settings. For instance, Rysman & Simcoe (2007) use these methods to show that patents associated with voluntary standard setting organizations have a significantly different citation age-profile than similar patents with no connection to compatibility standards. The flexible functional form is particularly useful in this setting, where the authors have little *a priori* information on how the age profiles might differ. More generally, our method provides a new approach to studying the rate of knowledge diffusion and obsolescence—an important issue that has arguably received less attention than knowledge spillovers across geographic or institutional boundaries. We discuss an instrumental variables extension of our method to the problem of measuring the age profile of productivity among research scientists (studied by Hall et al., 2007) using an instrument proposed by Jones (2007).

Finally, by exploiting the grant lag, our paper contributes to a growing literature studying sources of exogenous variation, and the grant lag in particular, in the patent process. For example, Gans, Hsu & Stern (2008) show that the likelihood of an innovation being cooperatively licensed increases on the patent's grant date. Murray & Stern (2007) link patents in the life sciences to associated publications in academic journals and show that journal citations decrease after the patent grant date.

The rest of the paper is as follows: Section 2 describes various age processes that might correspond to use of either the grant date or application date to define the age of a patent. Section 3 describes the data and explores the validity of using grant year instead of application year when defining the age of a patent. Section 4 shows the citation-age profile using this new definition of age, and presents results for different industries and time periods. Section 5 concludes.

# 2 Birth Dates and Citation Age

In papers by Jaffe & Trajtenberg (2002), Hall et al. (2002), and Caballero & Jaffe (2002), the citation age-profile is meant to reflect the process of an idea diffusing across the economy and then becoming obsolete as the technological frontier builds past it. However, in order to study the age-profile, we must choose a starting date for the "citation clock." Put differently, we must choose a suitable definition of a patent's age with respect to the citation-generating process. There are several possible candidates: the application date marks the day a successful application was filed; the grant date is the day when the USPTO issues a patent; and the priority date marks the beginning of an inventor's temporary monopoly.[5] In this section, we distinguish between different (admittedly stylized) processes that may affect the citation age-profile and discuss their relationship to either the application date or the grant date of a patent.

## 2.1 Age Processes that Start at the Application Date

In some instances, knowledge may diffuse quickly, so that it is used by others before a patent is published. Here we list some possible processes of transferring knowledge that might tend to start at the application date rather than the grant date:

[5]The priority date is the first date on which any related patent application is filed. In some cases, inventors can make changes to a pending application (resulting in a new application date) while keeping the original priority date. Thus, it is possible for the priority date to precede the application date listed on a patent. For a detailed discussion of the patent application process, see Popp et al. (2004).

1. Behavior of the assignee or inventor before patent publication: An idea may spread through the economy as soon as it is realized in several ways. The first is simply that in developing the idea, the inventor might present the idea at a conference or to colleagues to get feedback on the idea or know-how required. An idea may also spread via word-of-mouth or other informal means. It is possible that through employee mobility, the knowledge developed at one company can be used by another firm. Furthermore, knowledge can be transferred if an inventor exploits the innovation prior to the granting of a patent. This can occur through various avenues, including (but not limited to) preparing the means of production or increasing customer interest. In all of these cases the knowledge is spread to others who can either use the same technology for another application or use a different technology for the same application prior to the granting or publication of the patent.[6] To the extent that all of these processes are associated with the realization of an idea and not with the patent process, they would be better approximated by an age process beginning with the application date.

2. Economic or technological opportunity: Innovations are typically a response to a technological or economic need in society. As a result, several people may simultaneously research and develop a similar technology. The first assignee to apply for the patent will be cited by other applicants who later try and patent a similar innovation. The citation may be added during the patent application process by either the USPTO or the applicant due to a search of relevant patents. While these citations do not indicate a knowledge spill-over, they may be an important part of the overall citation age process. Indeed, it appears that more than 80 percent of the citations to a patent that has been applied for but not yet issued are added by the USPTO.[7] As these types of citations are associated with the realization of an idea rather than the diffusion process, they are better approximated as being associated with the application date.

3. Foreign publication: In most foreign countries, patent applications are published 18 months after the application date (rather than waiting until grant). This implies that U.S. patents that apply for foreign protection are published in a foreign country. To the extent that publication drives the age citation process, the fact that foreign publication

[6] These are the two generic relationships that Jaffe, Trajtenberg & Fogarty (2000) cite for the relationship between the cited patent and the citing patent.

[7] Thanks to Ajay Agrawal for providing data on pre-grant examiner added cites to address this question. For more evidence on the age profile of examiner added citations see Sampat & Hegde (2007).

is mechanically tied to the application date means that the application date better approximates the timing of this process in these cases. Since 2000, patents that applied for foreign protection are published 18 months after application in the U.S. as well (we further discuss this issue below).

## 2.2 Age Processes that Start at the Grant Date

Here, we describe some interactions that may lead to the age of a patent being better described starting from the grant or publication date:

4. Behavior of assignee and inventor after patent publication: The granting of a patent guarantees the protection of the idea and may lead the inventor to change behavior. For instance, the inventor may formally and informally present the technology. Also, the inventor may publish the innovation in a journal or industry publication. In addition, the inventor or assignee may be willing to begin commercializing the idea which means teaching the idea to suppliers, manufacturers and customers. This may lead to other parties exploiting the final product or some part of the production process in another application. Note that a granted patent provides protection back to the priority date, so the major issue that changes with grant is the resolution of any uncertainty that the applicant had about whether the patent would be granted. In practice, this issue may not be of much importance since there is little uncertainty as the great majority of applications are successful.

5. Domestic Publication: Granted patents are available from the USPTO in a searchable database. If an innovation is expanded or used because another inventor learned about the original innovation through a patent application then the diffusion process should start when the patent becomes public, or at the grant date of the application. Naturally, this process may be more relevant within certain technology classes (see Jaffe, Trajtenberg & Fogarty, 2000).

The relative importance of these different age processes is an empirical question. While we do not distinguish between all of them in this paper, we do provide evidence that the grant date is much more important than the application date in determining age profiles (in Section 3. Note that because our analysis controls for cohort and calendar effects (and because age since application is collinear with cohort and calendar effects), we implicitly control for age since application. However, we cannot separate the effect of age since application from the cohort and calendar effects.

One concern for the applicability of our method is that the American Inventors Protection Act of 1999 made the application date relatively more important. In particular, since 2000, U.S. patent applications associated with international families of applications (in the sense that they have the same priority date) are published 18 months after application rather than at the grant date. Johnson & Popp (2003) suggest that the impact may not be very large. They estimate that only a third of U.S. patents apply for international protection and so the majority are still under the old rule of publication at the grant date. In addition, U.S. patents filing for international protection have always had to be filed within 12 months of the U.S. application in order to maintain the U.S. priority date internationally, after which they would be subject to the standard foreign requirement of publication at 18 months after application. That is, patents filing for international protection have had their maximum time until publication reduced from 30 months (18+12) to 18 under the new rule. Since most patents are granted within 30 months, the real reduction is usually less than 12 months.[8]

# 3 The Significance of Grant Dates

This section discusses the data set, examines the cross-sectional and time-series variation in grant lags, and illustrates the significance of the grant date relative to the application date for determining citation age profiles.

## 3.1 Data

We require a data set on patents. An NBER project has collected a selection of information on every recent patent and made it easily available for purposes of research. Jaffe & Trajtenberg (2002) provides a comprehensive overview of these data, including a discussion of the citation-age profile. [9] Note that it is possible to download each patent document from the USPTO, so in theory it would be possible to collect patents and other information not included in

[8]Determining which patents are part of international families can be difficult and is not a part of the NBER patent data base so we do not pursue these issues here, although it appears to be an interesting issue for further research. In unreported results, we find little change when only using patent data from before 2000.

[9]We actually use a version of the NBER data set that has been updated to contain patents issued between 1999 and 2002. These data are freely available through Bronwyn Hall's web site: http:elsa.berkeley.edu˜bhhallbhdata.html.

the NBER data set. However, it is a laborious process. We focus on the NBER data patent database since it is easily available and widely used.

In the NBER patent database, we observe each patent's identification number, application year, grant date, technology class and all patents that are cited. Hence, we can find the set of all patents that cited a particular patent, and the dates of those citations, by searching the database appropriately. We observe several other variables that we do not focus on here, such as the assignee's name and the country of the patent owner. Note that although we observe the exact day of the patent grant, we observe only the year of the application, so we use just the years of the application and grant as the relevant data in this paper. While it might be possible to improve on this by going to the USPTO web site or by purchasing access to proprietary data sources, we prefer to work with data that is easily accessible to other researchers. The data set has information on all patents granted from 1963 to 2002 and all patent citations by patents granted from 1975 to 2002. The data set includes approximately 3.4 million patents and over 22 million citations.

In this study, we are interested in estimating the citation generation process. We observe all citations made by patents granted in 1975 or later, so we drop any patents that apply before 1975. (We drop based on the application date instead of the grant date because if we dropped patents that were granted before 1975, we would observe only a selection of patents that applied before 1975.) We also eliminate patents where the patent review process took longer than five years, as that would be very unusual. Doing so eliminates an additional 21,274 patents. The final data is composed of 25,217,424 observations. The number of patent applications that end in grant remains fairly steady around 67,500 from 1975 to 1984, when it begins an abrupt almost linear climb to 141,478 in 1995. The trend may well continue but is difficult to see at this point because we do not observe the full set of granted patents due to truncation of the data set.

## 3.2 Variation in Grant Lags

During the period of study, patents took an average of 2 years to be granted. Close to 80 percent of patents were granted within two years of application date, with 60 percent issued in the second year. Figure 1 exhibits the distribution of the lag between the application year and grant year of a patent in the full sample.

In addition to the variance of the lag in the cross section, there have been changes in the workings at the USPTO that have caused the grant lag to change over time. Figure 2 graphs the mean lag by patent application-year during the period 1975-1998. (We stop at 1998 to minimize the role of truncation in the graph.) We observe an increase in the lag for patents that apply after 1976. This increase is believed to be due to budget cuts at the USPTO in 1979 that led to fewer reviews being undertaken in that year, and leading to an increased number of patents queuing for review, particularly in later years. [10] This resulted in a longer grant lag on average during the early 1980s (see Popp et al., 2004). A second increase in mean pendency lags began in the 1990's. This slowdown is often attributed to an increase in the number of patent applications submitted to the USPTO, and a resulting backlog in the review process.

## 3.3 Setting the Citation Clock

For the grant lag to be a useful exogenous variable for identifying age profiles, it must have two features. First, the grant date must influence the start of the age process. And second, grant dates should not otherwise affect the citation profile. In other words, we require that the citation profile look the same for patents with different grant lags, except for a "shift" corresponding to the length of the lag. In this sub-section, we show that grant lags have a strong influence on what might be regarded as the start of the "citation clock" but not the shape of the age profile. It is still possible that the grant lag affects the level of citations as argued by Johnson & Popp (2003). While this creates problems for our estimator, we discuss how one could potentially instrument for this problem, and we argue in Section 4 that the effect is small in our data set.

To study the importance of the grant lag, we begin by defining age in the traditional way—the difference between the application year of the cited patent and the application year of the citing patent (i.e. the citing-year or calendar year). In this environment, we cannot estimate a separate set of age, year and cohort effects because of the identification problem. However, we can estimate how one age profile differs from another.

Let $C_{it}$ be the number of forward citations received by patent $i$ in year $t$. Let the parameters $\alpha_b$ and $\alpha_t$ represent vectors associated with dummy variables for the application year of patent $i$ and the application years ($t$) of patents that cite $i$. Let the parameters $\alpha_{al}$ represent vectors associated with dummy variables for age for each grant lag. That is, we estimate age profiles for each grant lag with separate sets of dummy variables. We exclude parameters $\alpha_{a0}$ for

[10] In 1979 there were only 48,854 patents granted, despite 66,102 patents being granted in 1978 and 61,819 patents being granted in 1980. In 1981 there were even more patents granted, 65,771 before a gradual decline until 1984.

patents granted in the year of application (grant lag of 0) as they would not be separately identified from $\alpha_b$ and $\alpha_t$. Our estimation function is

$$C_{it} = f_{it}(\alpha_b, \alpha_t, \alpha_{al})$$

where $a = t - b$.

In practice, we specify $f_{it}()$ to be the Poisson model and estimate via maximum likelihood. While the negative binomial model is popular for its less restrictive assumption on the dispersion of citations, Wooldridge (2002) argues that the Poisson model is preferable because, unlike the negative binomial, the Poisson model provides a consistent estimate of the conditional mean function even if the dispersion is misspecified, and because the impact of misspecification on standard errors can be addressed by using robust (White) standard errors. We use robust standard errors in all that follows. Formally, let $b_i$ and $g_i$ represent the application (birth) year and grant year of patent $i$. Normalize time to start at 1 rather than 1975 and let $T$ be the last year in the data set so $t \in \{1, 2, \ldots, T\}$. Let $I\{x = y\}$ be the indicator function that equals 1 when $x = y$ and 0 otherwise. We solve:

$$\begin{aligned}
\max_{\alpha_b,\alpha_t,\alpha_{al}} \log L &= \sum_{i=1}^{n} \sum_{\hat{t}=b_i}^{T} C_{i\hat{t}} \ln\left(m_{i\hat{t}}(\alpha_b, \alpha_t, \alpha_{al})\right) - m_{i\hat{t}}(\alpha_b, \alpha_t, \alpha_{al}) \\
\texttt{where} & \\
m_{i\hat{t}}(\alpha_b, \alpha_t, \alpha_{al}) &= \sum_{b=1}^{T} \alpha_b I\{b = b_i\} + \sum_{t=1}^{T} \alpha_t I\{t = \hat{t}\} + \sum_{a=1}^{T} \sum_{l=1}^{5} \alpha_{al} I\{a = \hat{t} - b_i\} I\{l = g_i - b_i\}.
\end{aligned}$$

In this regression, the coefficients $\alpha_{al}$ capture the difference in citation rates between patents with a grant lag of $l$ years and a grant lag of 0 years, in the $a^{th}$ year after application. We use $\hat{t}$ to refer to the time period associated with a particular observation, and $t$ to index the vector of citing-year parameters $\alpha_t$.

It is easier to understand our results from a graph than from a table of parameter estimates. Figure 3 graphs the expected proportion of cites received in each year since application for patents with different grant lags. Recall that our identification strategy requires that these age profiles look the same across grant lags except for a shift corresponding to the duration of the lag. This is precisely what we observe in Figure 3. If the application year was much more important than the grant year, all of the lines would overlap. Instead, we see that each distribution is shifted back approximately one year—suggesting the importance of the grant date in determining the age profile. Citations seem to peak about two years after the grant year and then slowly decline. Strikingly, as the grant lag increases by one year the peak is reached one year later.[11]

While Figure 3 shows that the grant date is important, we also require that the shape of the age profile is similar across grant lags. This is difficult to see in Figure 3 because all of the profiles continue until seventeen years after application. If the citation age process truly begins at the grant date, patents with longer lags will have truncated lives, which should increase the height of the line in Figure 3.

To correct for this problem, we use the estimates that generated Figure 3 to compute an adjusted probability distribution that begins at the grant date and ends after 12 years. Figure 4 shows the expected proportion of citations received by patents with different grant lags in the $a^{th}$ year since application. Once again, if the grant lag exogenously determines the start of the citation clock, we should see identically shaped lines shifted over by one year for each lag. This is much closer to what we observe than the series of overlapping lines that would be produced if the age process began at application.

Each age profile in Figure 4 rises rapidly until it reaches its peak approximately two years after the grant year, when a patent receives between 8 and 9 percent of its "lifetime" citations. The shape and the peak of the density function is very similar for all grant lags—but this similarity is especially striking for patents that have a lag of two years, three years and four years. For patents with a grant lag of one year, the density peaks at a higher proportion of citations and then decreases more rapidly. Also, patents with a five year lag between application and grant have a citation-age density that peaks at a lower level of citations and then declines significantly slower than the other profiles. Nevertheless, the overall picture is one of very similar distributions shifted over by one year. This pattern provides support for the relative importance of the grant date in determining the citation profile.

These preliminary results corroborate earlier work by Johnson & Popp (2003), who also test whether the grant date is important in the diffusion process. Their approach is to estimate separate age profiles for patents with different grant lags where age starts at grant. If the application date was important, they would find that patents with longer lags had faster declines in citations because these patents would be further into their age profile. Instead, Johnson and Popp find that patents have similar rates of decline in their citation propensity,

[11]The peak of the citation-age profile for patents with a one year lag occurs at about three years after application year; it occurs at about five years after grant date for patents with a two year lag. After a two year lag, as the lag increases by one year, the peak occurs one year later.

with both short and long lags. That result suggests that grant date is the appropriate date at which aging begins. Their approach differs from ours in several respects. First, they define the grant lag from the priority date, not the application date. Second, they define the birth date of citing and cited patents based on the grant date, not the application date. They solve the resulting identification problem by using functional form assumptions on the cohort and age effects similar to Hall et al. (2002). Given these different approaches, our joint results provide strong evidence in favor of the grant date as the start of the age process.

# 4 Estimating the Citation Age Profile

This section looks at the citation-age profile using the grant year of the cited patent to define the age of a patent. First we present our results assuming that the grant lag is exogenous to the citation process. Then we discuss the case in which patents with higher grant lags receive more citations, as suggested by Johnson & Popp (2003). Finally, we examine the age profile of patents from different technological fields, and from different time periods.

## 4.1 Main Results

We model the number of citations that patent $i$ receives in period $t$ as:

$$C_{it} = f_{it}(\alpha_b, \alpha_t, \alpha_a, \alpha_c) \tag{1}$$

Since $C_{it}$ is a count variable that frequently takes a value of zero, we specify $f_{it}$ to be a Poisson function and estimate via maximum likelihood. Let $b_i$, $g_i$ and $c_i$ be the application year (or birth year), grant year and technology class of patent $i$. Note that we define age as a function of grant, i.e. $a_{it} = t - g_i$. We solve:

$$\begin{aligned} \max_{\alpha_b,\alpha_t,\alpha_a,\alpha_c} \log L &= \sum_{i=1}^{n} \sum_{\hat{t}=b_i}^{T} C_{i\hat{t}} \ln\left(m_{i\hat{t}}(\alpha_b,\alpha_t,\alpha_a,\alpha_c)\right) - m_{i\hat{t}}(\alpha_b,\alpha_t,\alpha_a,\alpha_c) \\ &\texttt{where} \\ m_{i\hat{t}}(\alpha_b,\alpha_t,\alpha_a,\alpha_c) &= \sum_{b=1}^{T} \alpha_b I\{b=b_i\} + \sum_{t=1}^{T} \alpha_t I\{t=\hat{t}\} + \sum_{c=1}^{6} \alpha_c I\{c=c_i\} + \sum_{a=0}^{17} \alpha_a I\{a=\hat{t}-g_i\}. \end{aligned}$$

Unlike previous work in this area, we are able to consider a highly-flexible estimation equation. Our specification of the age profile function and the other explanatory variables could be considered non-parametric in the sense that all of our explanatory variables take on discrete values and we include dummy variables for each value of each variable, but we do this within the parametric form of a Poisson model. The Poisson model is useful because it provides a simple interpretation of the parameters but in fact, the specific functional form matters very little. Because we use all dummy variables, the post-estimation predicted value for any given set of covariates, $E[C|b, t, a, c]$, will be the same regardless of whether we specify a Poisson model, the negative binomial model, or a linear model estimated by OLS. [12]

The interpretation of the parameters is straightforward. The vector of parameters $\alpha_b$ measure the extent to which patent $i$ that applies in year $b_i$ is likely to be cited compared to patents applied for in the base year, 1975. Similarly $\alpha_t$ measures the extent to which a patent is likely to be cited in year $t$ relative to the base year of 1975. The parameters for the age dummies $\alpha_a$ measure the likelihood of citation at age $a_{it}$ relative to age zero (the grant year). The parameters $\alpha_c$ capture industry differences for industry $c_i$. We categorize patents as in Hall et al. (2002) so that the six potential categories for a patent application are: Drugs and Medical, Communications and Computers, Chemicals excluding Drugs, Electronics, Optics and Nuclear, Mechanical, and Other. The base category are patents that fall into the Other category.

Note that our estimation method is slightly different than Hall et al. (2002). For a dependent variable, they take the mean number of citations for all patents in a given cell—defined by $b$, $t$ and $c$ (In their case, $b$ and $t$ also define $a$). That is, their observation is a cohort and the dependent variable is continuous. We take each patent-year as a separate observation and our dependent variable is a citation count. Hence, they use non-linear least squares where we use a count data model. [13]

In Table 1 we report the results from the Poisson regression of Equation 1. The first set of coefficients show the effect of the application year of the cited patent, the cohort effect. These effects are decreasing throughout the sample period. Thus, after conditioning on age and citing-year, we find that citations received by an average patent have been declining over

[12] Angrist (2001) discusses this point at length. He argues in favor using linear models in the context of limited dependent variables, particularly when explanatory variables can be cast as sets of dummy variables.

[13] In practice, we store our data as a list of each of the realized counts for each cohort. For instance, there are 2,545 observations with application year 1995, grant year 1998, age 1 and industry 1 with 0 citations. We store each cohort as a single observation keeping track of count 2,545, which saves greatly on space and estimation time. A drawback is that we cannot compute statistics that rely on patent-level information, such as standard errors that are clustered at the level of the patent. Presumably, this matters little in this study with so many observations. Rysman & Simcoe (2007) use a smaller, more focused data set and so use patent-level clustering and even patent-level fixed effects

time. This finding presumably reflects the increasing number of patent grants. A standard interpretation is that patent quality has declined — either because of changes at the USPTO, or because inventors are patenting marginal ideas — or that recent invention cohorts have lower average "fertility" (i.e. generate fewer follow-on innovations).

The second set of coefficients are from the dummies for each citing year, the calendar year effect. These parameters tend to remain relatively stable with a few increases and decreases in the late 1970s and early 1980s, but then they monotonically increase. This pattern may reflect an increase in the number of patents granted per year over time. Also, previous studies that show that recent patents contain a greater number of citations. One explanation is the increased availability of computerized search methods.

As in Hall et al. (2002), the technological field estimates provide evidence that patents in the Computers and Communications field are the most highly cited. Patents in the Drugs and Medical and Electricity subsequently follow with respect to citations received. The category of patents in the Chemical field comes in next and the coefficients for category show that patents in the Mechanical field receive the fewest number of citations, even fewer than those in the Other category.

The last set of coefficients are for the various ages of the patent. The coefficients on age increase for the first year and then steadily decrease. After four years the patent is less likely to be cited than it was in the year it was granted. The results are similar when we only use pre-2000 data, which predates the American Inventor's Protection Act.

The solid line in Figure 5 shows the citation-age relationship graphically. The profile shows the expected number of citations received by patents for each age. As expected from the above coefficients, the profile reaches a peak at one year after the grant date and then slowly tapers off.[14]

It is important to note that defining "age" based on the grant year of the cited patent leads to the possibility of negative age values.[15] In the NBER data, 3.68% of all citations occur before the grant date of the cited patent. We dropped any observation with a negative age in our previous regressions. To check our results, we re-estimate including ages -2 and -1 (using an age of -2 as the omitted dummy variable). The resulting profile is shown in Figure 5 as the dashed line. The main results are confirmed. The profile is highly peaked and the peak again

[14] Until 1995 patents expired after 17 years from grant date but after the Uruguay Round of GATT, patent protection in the US extends 20 years from filing date.

[15] For example, suppose Patent A is applied for in 1981 and granted in 1983. Patent B applies in 1982 and is granted some time after 1983. Suppose the patent officer reviewing B adds A to B's list of citations. In our formulation, A receives a citation at age negative one, since B applies in 1982 and A is granted in 1983.

occurs at one year after grant or roughly three years after the application year.

The non-linear estimation of Jaffe & Trajtenberg (2002) using the grant date to define age finds a similar high early peak with a long tail. They find that the peak citation rate occurs four years after application. Our result that the peak occurs one year after grant year is consistent with a slightly earlier peak—approximately three years after the patent application. The height of the predicted citation frequency that Jaffe and Trajtenberg find using a non-linear estimation is similar to what we find here; at the peak the typical patent in our analysis receives just over one citation, while Jaffe and Trajtenberg suggest the average patent receives about 0.9 citations at the peak. Further comparisons (described in Section 4.3) suggest that these differences stem largely from the use of different data sets rather than from the alternative econometric approaches. From a substantive viewpoint, these results are broadly consistent with Cabalerro and Jaffe's conclusion that while diffusion is nearly "instantaneous" (2002, p. 146) the knowledge embedded in a patent remains useful over rather long periods of time.

## 4.2 Endogeneity of the lag

The key to our identification strategy is the assumption that grant lags can be excluded from the citation generating function. To gain intuition, consider a linear model of citations:

$$C_{it} = \alpha_0 + \alpha_b b_i + \alpha_t t + \alpha_a a_{it} + \alpha_c c_i + \varepsilon_{it}$$

where, as before, $b_i$ is the birth (application) year, $t$ is the calendar year, $c_i$ is the technology class, and $a_{it} = t - g_i$ is the patent's age. Naturally, a key assumption for estimation is that $E[\varepsilon_{it}|b_i, t, a_{it}, c_i] = 0$. An important variable that might be contained in $\varepsilon_{it}$ is the pendency lag $l_i$. This lag must be excluded from this equation as it would be linearly dependent on the other variables: $l_i = t - a_{it} - b_i$. For the same reason, it is obvious that the grant lag is not uncorrelated with the explanatory variables; it is mechanically correlated with them by construction. Hence, it is crucial that the grant lag be uncorrelated with citations in order for this regression to be consistent, formally $E[\varepsilon_{it}|l_i] = 0$ or equivalently: $E[C_{it}|b_i, t, a_{it}, c_i] = E[C_{it}|b_i, t, a_{it}, c_i, l_i]$.

What if the lag is correlated with citations? For instance, Johnson & Popp (2003) have suggested that more important patents take longer to process so that patents with a longer grant lag (measured from the priority date) receive more citations. The phenomenon creates an endogeneity problem: patents with long lags have low ages (by construction) and high citations, so we underestimate the age parameters. The work of Regibeau & Rockett (2004) implies that this problem may be particularly severe since they show (in the context of genetically modified foods) that patents take longer to approve when they come from a relatively new technology, which is often also when we observe the most important patents.

One approach to addressing this problem would be to use instrumental variables. For instance, one could predict age ($a_{it}$) in a first stage regression as a function of birth year ($b_i$), calendar time ($t$) and an instrumental variable that predicts the grant lag but is otherwise uncorrelated with citations ($C_{it}$). In the context of patents, patent examiner fixed effects might be a useful instrument (for a discussion of patent examiners, see Cockburn et al., 2002). Using instrumental variables in count data models presents some problems, which could be addressed with appropriate non-linear modeling. [16]

In Section 4.5, we describe how one might use an instrumental variable approach with our method. Here, we take a different approach, and argue that the size of any endogeneity bias is small in our application. To examine the size of the bias produced by our estimation procedure, we re-estimate Equation 1 on two different samples. The first sample contains a random selection of patents with a grant lag of less than or equal to two years, and the second adds all patents with a grant lag of three years. If our estimates of the age profile are strongly biased due to the correlation between grant lag and citations, we would expect a large decrease in the age coefficients for the second regression.

For this exercise, we chose a random sample of patents with a zero to two year grant lag in order to reduce the size of the first sample. When we use the full sample, there is virtually no change in the results from adding patents with a grant lag of three years since roughly 80 percent of all patents have a grant lag less than three years. By reducing the sample size of the patents with grant lags less than three years, we create a strong bias in favor of observing larger changes. However, as the selection of data sets leads us to over represent some years, the parameters differ a bit from Table 3.

The results of this bias-testing exercise appear in Table 2. As expected, the age coefficients decline when patents with a three-year grant lag are added to the estimation sample. The differences are initially quite small, but they increase with age. For example, the Age 1 coefficient for patents with a lag of two years or less is 0.093 and for patents with lags of

[16]In particular, a standard approach for addressing endogeneity in a non-linear model is to explicitly model the correlation between unobservable terms in the "first" and "second" stage, and estimate the two equations simultaneously, for instance via full information maximum likelihood. However, the Poisson model does not explicitly state the distribution of the unobservable term so it is difficult to allow for correlation between the unobservable term in a Poisson model and the unobservable term in some other equation. See Mullahy (1996),Windmeijer & Santos-Silva (1997) and Miravete (2008) for some progress on this issue.

3 years or less is 0.046 — a difference of 0.047. Since these coefficients come from a Poisson regression, the difference can be interpreted as a percentage change in the baseline citation rate. Three years after the grant date, the relevant coefficients are -0.086 and -0.182 – a difference of 0.096 – and at age 6, the difference is 0.146. However, it is important to keep in mind that patents receive very few citations at these high ages, so a large difference in coefficients (which corresponds to a percentage difference) still means a relatively small difference in expected citations.

To further explore this issue, the first panel in Figure 6 graphs the expected number of citations at each age for the results from each sample.[17] The difference in predictions appears rather small. The difference in citations at age 1 is 0.0055 citations, the maximum is 0.0091 citations at age 4 and the difference declines slowly to 0.0077 at age 10 and 0.0062 at age 15.

The second panel in Figure 6 examines the probability density functions for each age profile (i.e. normalizes the fitted values by expected cumulative citations). The largest differences appear during the first few years after grant. While the age profile for the smaller sample is more spread out – suggesting that our approach will overstate the mean citation age – this difference does not appear to be unreasonably large.

Overall, this exercise suggests that our identification strategy estimates a lower bound for the true age coefficients. While our analysis suggests that the bias is relatively small in large samples, we cannot rule out the possibility that it would be important in the context of a particular application.

## 4.3 Differences across industries

In this section, we consider how age profiles differ across technological fields, following Hall et al. (2002). Specifically, we re-estimate Equation 1 including a full set of interactions between age and technological field dummies. This allows patent citations in each technological field to follow a different age profile.

Figure 7 shows the citation profile for patents by industry. We find very intuitive results. The pharmaceutical industry ("Drugs"), for which patent protection is a crucial issue, has the longest active citation life. The communication and computer industry, for which we believe technological turnover is very fast, has the shortest active citation life. The electrical and electronics industry is almost as fast. The chemicals and manufacturing industries fall

[17]For both computations, the expected number of citations is computed with all dummy variables (but for those associated with age) set to zero.

between drugs and the computers/electronics set.

We now compare these results to those in Hall et al. (2002). The difference we wish to highlight is the comparison between our use of a non-parametric age profile function (under the assumption that aging starts at grant) and their use of a double exponential parametric functional form. However, their implementation differs from our in several respects, so it is useful to consider which one drives the differences. First, they use the grant year of the cited and citing patent for the birth year and citing year respectively. Second, they group birth year into 5-year bins for creating dummies, rather than treating each birth year with a separate dummy. Third, they aggregate patents into cohorts defined by grant year and category, and look at the average number of citations by citation year. Thus, their dependent variable is continuous and they use non-linear least squares rather than the Poisson model that we use. Finally, and we find most importantly, our data sets do not match. They use patent data going back to 1965, whereas we start our data set in 1975.

To compare the two methods, we begin by estimating their model on our data, and then estimate models successively closer to our own to determine which modeling changes produce significant changes in the results. In practice, we do not find exactly the same parameter values as they do but we consistently find the same basic hump shape in the age profile (Hall et al., 2002, Figure 20).

Since the two methods generate a different set of parameter vectors, we focus on differences in the mean citation age. We compute the mean citation age by calculating $p_a$, the probability of a citation arriving in each of the first 15 years after grant (including age 0), and then computing $\sum_{a=0}^{15} p_a a$. [18] Hall et al. (2002) focus on the mode (i.e. the "peak" in the citation age profile) rather than the mean. We examined the mode and found similar results, but it is less interesting in our model because the mode is always an integer. Finally, because our data set uses an enormous number of observations, any differences are statistically significant and we focus on changes in the relevant statistics without reporting standard errors.[19]

Results appear in Table 3. The first column reports the modes of each distribution from the results of Hall et al. (2002). We simply copy this column from their book for comparison purposes. The second replicates their model in our data. That is, we aggregate our data to

[18]Probability weight above 15 is therefore assigned proportionally to all years rather than being summed into the final year. Also, each age is treated as a separate discrete bin. An alternative would be to integrate over age treating it as a continuous variable, but this would only be sensible for the specifications with the double exponential functional form for the age profile.

[19]Of course, it is possible to compute a standard error for the mean citation age using the delta method: see Rysman & Simcoe (2007) for one such application.

cohorts defined by grant year and cohort, compute the average number of citations in each age (where age is defined as the grant year of the citing patent minus the grant year of the cited patent). Birth year dummies represent five-year intervals and we use their double exponential functional form for the age profile.[20] We allow the two coefficients in the exponential functional form to differ by industry. Column 2 reports the mean of the age of a citation from this calculation.

Comparing columns 1 and 2 finds similar results. They find that the catch-all industry "other" has the longest active citation life, which we do not find. However, excluding "other", the pharmaceutical industry ("Drugs") experiences citations latest in life in both data sets. Hall et al. (2002) find that "Electrical" and "Manufacturing" are very close and have the shortest active citation lives, with "Computers" and "Chemicals" being grouped together having slightly longer citation lives, but still well short of drugs. This is similar to the result in our data, in which "Electrical", "Manufacturing", "Computers" and "Chemicals" have mean citation ages within a third of a year of each other. We do not report the mode that we find in our data but we can compute it. As before, it arrives somewhat earlier than that found in Hall et al. (2002).

The next column does not use patent cohorts, and instead treats each patent separately (that is, each patent-year is a separate observation). Other than that, we use their approach of grouping birth years for purposes of estimating birth year dummies, using grant years to define birth and citing year and using the two-parameter double-exponential age function with separate parameters for each industry. The fourth column repeats this exercise using the application year to define birth and citing year. Column 5 reports the results from the model we advocate in this paper, where we estimate a full set of birth year, citing year and age dummies and define age by the grant year of the cited patent and the application year of the citing patent. Results are remarkably similar across these specifications. "Drugs" has the latest mean citation date, followed by "other". Computers and electronics have the earliest mean citation ages. Chemicals and manufacturing form a middle group, between drugs and computers/electronics. The ordering of the six industries is identical across specifications 2 through 4, while 5 finds that "Computer" citations arrive slightly faster than "Electrical."

Note that the final column reports numbers somewhat less than the previous columns. The difference ranges from 1.5 to 1.9. That is because columns 2, 3 and 4 report age as the difference

[20]In particular, we let the expected number of citation be function of $f(a)$ where $f(a) = \exp(-\beta_1 a)(1 - \exp(-\beta_2 a))$. Hence, the log of the expectation is affected by $\ln(f(a))$. This set-up follows Caballero & Jaffe (2002).

between either grant years or application years, whereas column 5 uses the difference between the grant year of the cited patent and the application year of the citing patent. Recalling that patents average 1.5 to 2 years between application and grant, the results in column 5 are perfectly consistent with the results in columns 2, 3, and 4. However, to show that the inter-industry differences produced by these two methods are quite similar when applied to the same data set, the last two columns in Table 3 report differences relative to "Other" for models 2 and 5.

Overall, our technique finds the same results as the technique proposed by Hall et al. (2002). The differences between our results and theirs appear to be driven by data issues rather than methods. Hence, we take this as strong confirming evidence in favor of these results about the hump shaped citation pattern and the relative differences across industries. It is striking to see two different sets of assumptions arrive at such similar results. [21]

## 4.4 Differences over time

While differences in the age profile of citations across industries have been previously studied, there has been little direct work on how the age profile has changed over time. Given the large changes in the administration and judicial treatment of patents, as well as the evolving relationship between research and commercial activity, one might imagine that the nature of diffusion and obsolescence has changed. In fact, we find that the age profile of citations has remained remarkably stable over the last three decades.

In order to explore this issue, we specify a model similar to that in the previous subsection. We use a Poisson model to predict citations to a patent in a year with a full set of cohort and calendar year dummies. We also interact a full set of age dummies with 5-year cohort dummies, allowing the age profile to change across 5-year spans. Hence, we interact age with cohort rather than industry. We use patents with application years from 1975 to 1999 for this exercise (so there are five years in each cohort). As above, we date the cohort year with the cited patent's application year, the age process beginning with the cited patent's grant year and the age based on the citing patent's application year.

A graph of the probability density function of the age profile appears in Figure 8. The lines appear very similar, almost perfectly overlapping for the first three cohorts, and only slightly

[21] Interestingly, Hall et al. (2002) report modal citation ages for the case of one parameter in their age profile equation being held constant (the "diffusion" parameter) which match the ordering and grouping that our model finds exactly. (pg. 445)

different for the fourth cohort. The last cohort, 1995-1999, does appear somewhat different. Although the peak is the same as for the other cohorts, the downward slope after the peak appears less steep. While this change in the age-profile suggests that recent patents have been more productive at ages one to seven, this result may also be an artifact of the truncated set of data for this cohort. It will be interesting to revisit this result in several years.

One interesting difference appears in year zero, where the first two cohorts receive more citations than the next two. An advantage of our estimation approach is that it handles age zero (or even negative ages) as it does any other age. In contrast, the functional form used by Hall et al. (2002) cannot address citations at age zero since it uses the term $\ln(1-\exp(\beta_2 a))$. Hence, they normalize citations at age zero to zero for each industry. While our results suggest that this normalization is restrictive, citations in year zero are difficult to interpret. In particular, they are more likely to have been added by legal counsel or patent examiners, and thus may not reflect diffusion of the underlying ideas. For instance, while the overall incidence of examiner-added citations is 39.7% (Hegde and Sampat 2007, Table 1) this figure increases to 78% for citations received at ages less than or equal to zero (i.e. at or prior to grant).[22]

Table 4 presents mean citation ages for each grant-year cohort. We only use up to age seven, as that is all we have for the last cohort. Results are similar if we just use up to age five. The first row uses ages zero to seven, and shows a monotonic increase in the mean age from 2.83 to 2.97 over the first four cohorts. However, this may be misleading as much of the difference is due to changes in citations at age zero, which may not reflect the diffusion process we are interested in. The second row calculates the mean age using only ages 1 to 7, and we see that the change in the age of a citation is practically eliminated. The mean age ranges from 3.41 to 3.47 for the first four cohorts, although interestingly, the change is still monotonic in cohort. We conclude that the age profile has changed very little, if at all. However, for both calculations, we see that the mean age for the last cohort is substantially higher. While this result is potentially quite interesting, we prefer to wait for more data before drawing conclusions.

[22]Data on examiner added citations are only available for patents granted since 2001. Hegde and Sampat also find a sharp drop in the share of cites added by examiners between 1 and 4 years after application, which is broadly consistent with our findings on the importance of the grant date. We calculated the 78% figure using data on grant-lags and examiner-added citations that was generously provided by Ajay Agrawal.

## 4.5 Extensions

Since the age-year-cohort identification problem occurs in a variety of different settings, our method might be applied to many problems where there is a lag between "birth" and the start of the "aging" process. However, in many applications, researchers will likely be concerned with the assumption that these lags are exogeneous. Here, we provide a brief discussion of how one might use our method given an instrumental variable that produces exogenous variation in "age" conditional on birth cohort. To frame this discussion, we focus on the applications of Hall et al. (2007) and Jones (2007). Our goal is to show the wide applicability of our method and, as such, we do not provide any new empirical results.

Hall et al. (2007) study the life-cycle productivity of scientists using data on publications by 465 French physicists born between 1936 and 1960. They are interested in estimating how the expected number of publications changes with age (the age profile of publications) but also must control for calendar and cohort affects. They compute age as the literal age of the scientist, i.e. the calendar year minus the year in which the scientist was born, and use functional form assumptions on cohort and calendar affects to separate the resulting endogeneity problem.

If the physicists in Hall et al. (2007) do not begin publishing until their course work is completed, exogenous variation in the time to complete their degree would provide useful variation. That is, our approach will be useful if scientific age differs from physical age. This approach is taken up by Jones (2007), who studies the age at which scientists have a Nobel prize-winning idea, using World War I and World War II as exogenous shocks that lengthened the time needed to complete a PhD or equivalent degree. [23]

To apply our approach in this setting, a researcher would first compute each scientist's age-since-degree (the calendar year minus the final year of course work), and then estimate a regression that includes a separate set of dummy variables for calendar years, birth-year cohorts and age-since-degree. This approach requires three assumptions: 1) publication begins once course work is complete; 2) scientific output is not correlated with time spent in graduate school; and 3) conditional on output, the impact of age-since-degree on publication rates is not correlated with time spent in graduate school. These are the same assumptions we make above: assumption 3 says that the shape of the citation age-profile does not depend on "degree lags" while assumption 2 says that these lags are independent of a scientists' total output. While Jones (2007) argues in favor of assumption 1, the latter two assumptions might be debatable, in

[23] Jones (2007) also captures the age profile in a two-parameter model. He does not estimate a calendar or cohort effect. Presumably, this is because the number of Nobel Prizes is fixed at one per year, and hence does not vary with calendar or cohort years.

which case age-since-degree would be endogenous to the number of publications. For instance, assumption 2 could be problematic if someone who completes graduate school quickly also has high ability.

As a solution, we could predict age in a first stage regression with dummy variables that indicate whether an individual was of the appropriate age to have the completion of their studies delayed by World War II. Then, a second-stage regression using predicted age would consistently estimate the correct parameters. [24] If assumption 3 was wrong, then the entire shape of the age profile could be endogenous. That would be true if scientists are more productive when they are younger so delays push them into less productive periods. This is exactly the issue studied by Jones. In this case, the first stage might use the World War II dummy interacted with calendar time or birth year or both.

# 5 Conclusion

This paper provides a method for flexibly estimating the age profile of patent citations while controlling for cohort and citing-year effects, a problem that has vexed previous researchers. We rely on the exogeneity of the grant lag—the time between the application and grant of a patent. If the relevant date for the birth of a patent is the application year but the aging process does not start until the grant date, then the age profile can be identified without resorting to functional form assumptions on the age profile.

To justify our identification strategy, we show empirically that the grant date is more important than the application date in determining the citation age profile. This finding is relevant for any researcher that must choose the appropriate "birth date" for a patent, regardless of whether they are actually estimating a citation age-profile. We also explore a potential problem with our approach: correlation between the grant lag and citation levels. While there is some evidence of a positive correlation, we argue that it produces a relatively small downward bias on the age coefficients in our application.

When we compare our estimates of the citation age profile to those obtained using more parametric methods, we find strong similarities. In particular, both methods produce a hump-shaped profile that peaks early in the life of a patent and declines gradually (though the peak, or modal citation age, arrives somewhat earlier in our approach). In comparing across

[24] As noted in Footnote 16, literally using two-stage least squares would require a linear second stage, not the Poisson model we use in this paper. That might be appropriate if the dependent variable is continuous. Otherwise, some alternative estimation strategy must be employed.

technological fields, we find that our method corroborates the results produced from a previous method. Finally, we present some new results showing that the shape of the citation age profile was remarkably stable for patents granted between 1975 and 1994. While there is some evidence of a change in the age profile for patents in recent grant-year cohorts, more data is needed to properly test that hypothesis.

While this paper focused on estimating age effects in patent data, we conclude by noting that the age-year-cohort identification problem occurs in many different settings. Thus, given an instrument that produces exogenous variation in "age" conditional on birth cohort, our method might be applied to a wide variety of problems. For instance, one could analyze citations to academic journal articles using the lag between submission and publication. In that setting, editor fixed effects might represent a promising set of instruments. Overall, we believe that the basic idea of relying on exogenous delays in the aging process to separate cohort, calendar and age effects has many potential applications.

# References

Adams, J., Clemmons, J. R., & Stephan, P. (2006). How rapidly does science leak out? NBER Working Paper No. 11997.

Agrawal, A., Cockburn, I., & McHale, J. (2006). Gone but not forgotten: knowledge flows, labor mobility, and enduring social relationships. *Journal of Economic Geography*, *6*(5), 571–591.

Albert, M. B., Avery, D., Narin, F., & Mcallister, P. (1991). Direct validation of citation counts as indicators of industrially important patents. *Research Policy*, *20*(3), 251–259.

Alcacer, J. & Gittelman, M. (2006). Patent citations as a measure of knowledge flows: The influence of examiner citations. *Review of Economics and Statistics*, *88*(4), 774–779.

Angrist, J. (2001). Estimation of limited dependent variable models with dummy endogenous regressors: Simple strategies for empirical practice. *Journal of Economics and Business Statistics*, *19*, 2–28.

Bacchiocchi, E. & Montobbio, F. (2004). EPO vs. USPTO citation lags. CESPRI Working Paper n. 161.

Branstetter, L. & Ogura, Y. (2005). Is academic science driving a surge in industrial innovation? evidence from patent citations. NBER Working Paper No. 11561.

Caballero, R. J. & Jaffe, A. B. (2002). How high are the giants' shoulders: An empirical assessment of knowledge spillovers and creative destruction in a model of economic growth. In A. B. Jaffe & M. Trajtenberg (Eds.), *Patents, Citations and Innovations: A Window into the Knowledge Economy*. Cambridge: MIT Press.

Cockburn, I., Kortum, S., & Stern, S. (2002). Are all patent examiners equal? impact of patent examiner characteristics. NBER Working Paper No. 8980.

Galenson, D. W. & Weinberg, B. A. (2000). Age and the quality of work: The case of modern american painters. *Journal of Political Economy*, *108*, 761–777.

Gans, J. S., Hsu, D. H., & Stern, S. (2008). The impact of uncertain IP rights on the market for ideas: Evidence from patent grant delays. *Management Science*, *54*, 982–987.

Hall, B., Jaffe, A., & Trajtenberg, M. (2002). The NBER patent citations data file: Lessons, insights, and methodological tools. In A. Jaffe & M. Trajtenberg (Eds.), *Patents, Citations and Innovations: A Window on the Knowledge Economy*. Cambridge: MIT Press.

Hall, B., Jaffe, A., & Trajtenberg, M. (2005). Market value and patent citations. *RAND Journal of Economics*, *36*, 16–38.

Hall, B. H., Mairesse, J., & Turner, L. (2007). Identifying age, cohort and period effects in scientific research productivity: Discussion and illustration using simulated and actual data on french physicists. *Economics of Innovation and New Technology*, *16*, 159–177.

Harhoff, D., Narin, F., Scherer, F. M., & Vopel, K. (1999). Citation frequency and the value of patented inventions. *Review of Economics and Statistics*, *81*(3), 511–515.

Jaffe, A. & Trajtenberg, M. (2002). *Patents, Citations and Innovations: A Window on the Knowledge Economy.* Cambridge: MIT Press.

Jaffe, A., Trajtenberg, M., & Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *Quarterly Journal of Economics*, *108*, 577–598.

Jaffe, A. B. & Trajtenberg, M. (1996). Flows of knowledge from universities and federal laboratories: Modeling the flow of patent citations over time and across institutional and geographic boundaries. *Proceedings of the National Academy of Sciences of the United States of America*, *93*(23), 12671–12677.

Jaffe, A. B. & Trajtenberg, M. (1999). International knowledge flows: evidence from patent citations. *Economics of Innovation and New Technology*, *8*, 105–136.

Jaffe, A. B., Trajtenberg, M., & Fogarty, M. S. (2000). Knowledge spillovers and patent citations: Evidence from a survey of inventors. *American Economic Review*, *90*, 215–218.

Johnson, D. K. N. & Popp, D. (2003). Forced out of the closet: The impact of the American Inventors Protection Act on the timing of patent disclosure. *RAND Journal of Economics*, *34*, 96–112.

Jones, B. F. (2007). Age and great invention. Unpublished Manuscript, Kellogg School of Management.

Kortum, S. & Lerner, J. (1999). What is behind the recent surge in patenting? *Research Policy*, *28*(1), 1–22.

Marco, Alan, C. (2007). The dynamics of patent citations. *Economics Letters*, *94*, 290–296.

Miravete, E. J. (2008). Competing with menus of tariff options. Unpublished manuscript, University of Texas-Austin.

Mowery, D. C., Oxley, J. E., & Silverman, B. S. (1996). Strategic alliances and interfirm knowledge transfer. *Strategic Management Journal*, *17*, 77–91.

Mullahy, J. (1996). Instrumental-variable estimation of count data models: Applications to models of cigarette smoking behaviour. *Review of Economics and Statistics*, *11*, 586–593.

Murray, F. & Stern, S. (2007). Do formal intellectual property rights hinder the free flow of scientific knowledge? An empirical test of the anti-commons hypothesis. *Journal of Economic Behavior and Organization*, *63*(4), 648–687.

Popp, D., Juhl, T., & Johnson, D. K. N. (2004). Time in purgatory: Examining the grant lag for U.S. patent applications. *Topics in Economic Analysis and Policy*, *4*, 1–43.

Regibeau, P. & Rockett, K. (2004). Are more important patents approved more slowly and should they be? Unpublished Manuscript, University of Essex.

Rysman, M. & Simcoe, T. (2007). Patents and the performance of voluntary standard setting organizations. forthcoming, Management Science.

Sampat, B. & Hegde, D. (2007). Examiner citations, applicant citations, and the private value of patents. Unpublished Manuscript.

Thompson, P. & Fox-Kean, M. (2005). Patent citations and the geography of knowledge spillovers: A reassessment. *American Economic Review*, *95*, 450–460.

Trajtenberg, M. (1990). A penny for your quotes: Patent citations and the value of innovations. *RAND Journal of Economics*, *21*, 172–187.

Windmeijer, F. & Santos-Silva, J. (1997). Endogeneity in count data models: An application to demand for health care. *Journal of Applied Econometrics*, *12*, 281–294.

Wooldridge, J. M. (2002). *Econometric Analysis of Cross Section and Panel Data*. MIT Press.

Table 1: Main Results

| | Application Year | | Citing Year | | Age Since Grant | | |
|---|---|---|---|---|---|---|---|
| | Coeff. | S.E. | Coeff. | S.E. | Age | Coeff. | S.E. |
| 1976 | 0.005 | (0.003) | 0.308 | (0.100) | 1 | 0.165 | (0.002) |
| 1977 | -0.016 | (0.003) | 0.410 | (0.099) | 2 | 0.160 | (0.002) |
| 1978 | -0.048 | (0.004) | 0.414 | (0.099) | 3 | 0.085 | (0.003) |
| 1979 | -0.062 | (0.004) | 0.447 | (0.099) | 4 | -0.012 | (0.003) |
| 1980 | -0.075 | (0.005) | 0.430 | (0.099) | 5 | -0.121 | (0.004) |
| 1981 | -0.092 | (0.005) | 0.412 | (0.099) | 6 | -0.229 | (0.004) |
| 1982 | -0.106 | (0.006) | 0.445 | (0.099) | 7 | -0.332 | (0.005) |
| 1983 | -0.096 | (0.006) | 0.426 | (0.099) | 8 | -0.432 | (0.005) |
| 1984 | -0.090 | (0.006) | 0.528 | (0.099) | 9 | -0.529 | (0.006) |
| 1985 | -0.096 | (0.007) | 0.609 | (0.099) | 10 | -0.618 | (0.007) |
| 1986 | -0.092 | (0.007) | 0.692 | (0.099) | 11 | -0.704 | (0.007) |
| 1987 | -0.123 | (0.008) | 0.795 | (0.099) | 12 | -0.784 | (0.008) |
| 1988 | -0.157 | (0.008) | 0.899 | (0.099) | 13 | -0.867 | (0.009) |
| 1989 | -0.189 | (0.009) | 0.942 | (0.099) | 14 | -0.940 | (0.009) |
| 1990 | -0.218 | (0.010) | 0.984 | (0.099) | 15 | -1.015 | (0.010) |
| 1991 | -0.239 | (0.010) | 1.013 | (0.099) | 16 | -1.093 | (0.010) |
| 1992 | -0.264 | (0.011) | 1.108 | (0.099) | 17 | -1.168 | (0.011) |
| 1993 | -0.299 | (0.011) | 1.212 | (0.099) | | | |
| 1994 | -0.337 | (0.012) | 1.370 | (0.099) | Technological Field | | |
| 1995 | -0.425 | (0.013) | 1.625 | (0.099) | | | |
| 1996 | -0.478 | (0.013) | 1.585 | (0.099) | Chemicals | | 0.023 |
| 1997 | -0.590 | (0.014) | 1.755 | (0.099) | Computers | | 0.647 |
| 1998 | -0.647 | (0.015) | 1.674 | (0.099) | Drugs and Medical | | 0.507 |
| 1999 | -0.627 | (0.016) | 1.594 | (0.100) | Electricity | | 0.232 |
| 2000 | -0.618 | (0.020) | 1.258 | (0.100) | Mechanical | | -0.054 |
| 2001 | -0.908 | (0.061) | 0.203 | (0.100) | Constant | | -1.524 |
| 2002 | | | -2.974 | (0.100) | | | |

Specification: Poisson; Observations: 25,217,424

Table 2: Bias Evaluation

| Age | Grant Lags 0,1,2 | | Grant Lags 0,1,2,3 | |
|---|---|---|---|---|
| | Coeff. | S.E. | Coeff. | S.E. |
| 1 | 0.093 | (0.009) | 0.046 | (0.006) |
| 2 | 0.042 | (0.011) | -0.033 | (0.007) |
| 3 | -0.086 | (0.014) | -0.182 | (0.008) |
| 4 | -0.228 | (0.018) | -0.347 | (0.010) |
| 5 | -0.388 | (0.021) | -0.525 | (0.011) |
| 6 | -0.541 | (0.025) | -0.687 | (0.013) |
| 7 | -0.669 | (0.029) | -0.843 | (0.015) |
| 8 | -0.791 | (0.033) | -0.990 | (0.017) |
| 9 | -0.923 | (0.036) | -1.153 | (0.019) |
| 10 | -1.053 | (0.040) | -1.304 | (0.021) |
| 11 | -1.149 | (0.044) | -1.440 | (0.023) |
| 12 | -1.271 | (0.048) | -1.584 | (0.025) |
| 13 | -1.373 | (0.052) | -1.693 | (0.027) |
| 14 | -1.490 | (0.056) | -1.817 | (0.029) |
| 15 | -1.617 | (0.060) | -1.958 | (0.031) |
| 16 | -1.763 | (0.064) | -2.102 | (0.033) |
| 17 | -1.858 | (0.068) | -2.210 | (0.035) |
| Obs. | 3,125,261 | | 5,711,815 | |

Cohort and citing-year effects not reported

| | (1) | (2) | (3) | (4) | (5) | Relative to "Other" | |
|---|---|---|---|---|---|---|---|
| Model | JHT Mode | JHT Replication | Poisson Grant-year | Poisson App-year | Our Model | Col 2 | Col 5 |
| Chemicals | 3.82 | 7.47 | 7.59 | 7.56 | 5.86 | -0.17 | -0.23 |
| Communication & Computers | 3.83 | 7.28 | 7.32 | 7.28 | 5.43 | -0.36 | -0.66 |
| Drugs | 4.75 | 8.04 | 8.19 | 8.14 | 6.32 | 0.40 | 0.22 |
| Electrical & Electronics | 3.27 | 7.24 | 7.28 | 7.23 | 5.49 | -0.39 | -0.61 |
| Manufacturing | 3.26 | 7.36 | 7.42 | 7.41 | 5.72 | -0.27 | -0.38 |
| Other | 5.55 | 7.63 | 7.73 | 7.72 | 6.10 | 0.00 | 0.00 |

Table 3: Mean age of a citation by industry

| Cohort | 75-99 | 80-84 | 85-89 | 90-94 | 95-99 |
|---|---|---|---|---|---|
| Ages 0-7 | 2.83 | 2.83 | 2.90 | 2.97 | 3.29 |
| Ages 1-7 | 3.41 | 3.41 | 3.42 | 3.47 | 3.78 |

Table 4: Mean cite-age by grant-year cohort

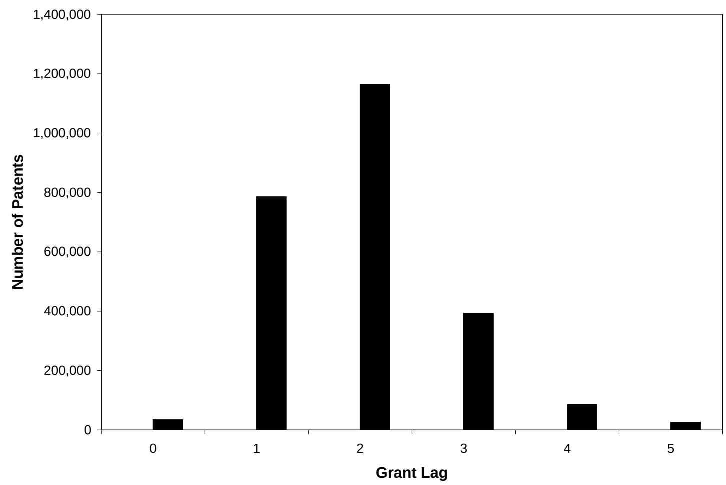


Figure 1: Frequency of grant lag

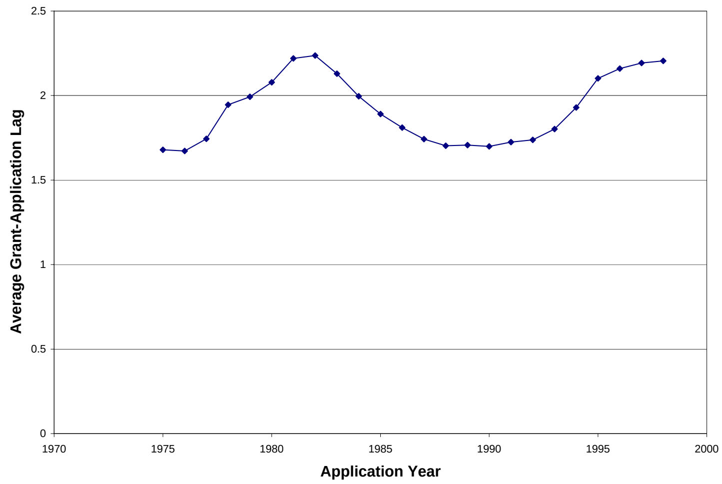


Figure 2: Average grant lag over time

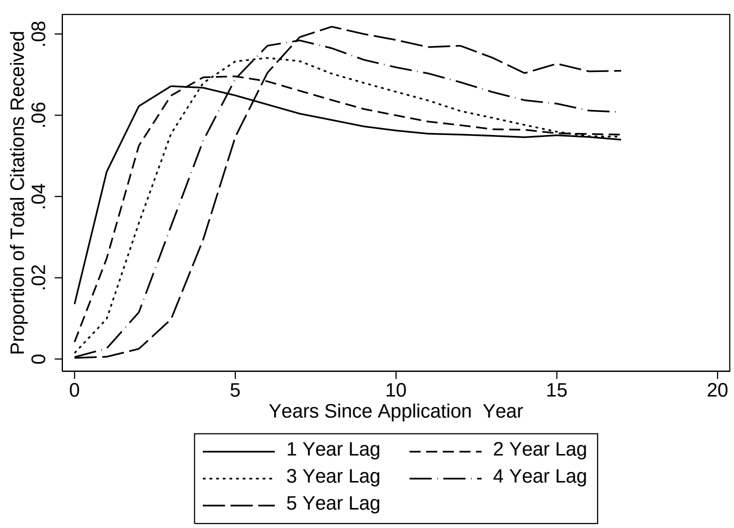


Figure 3: Effects of age (since application) for different grant lags

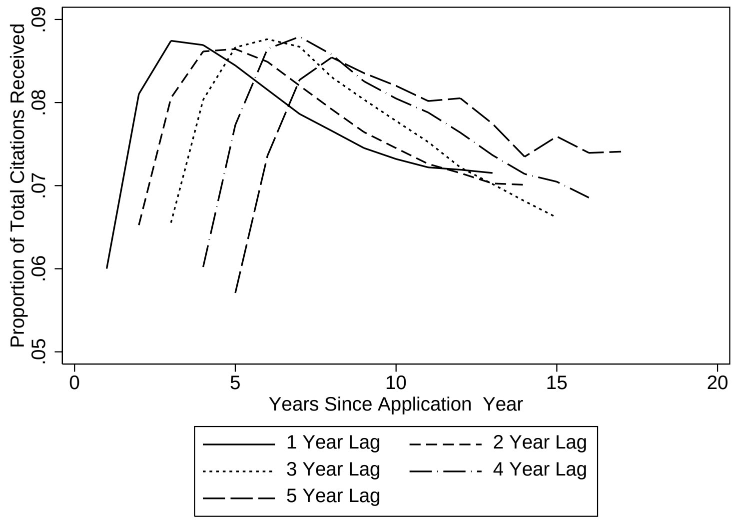


Figure 4: Effects of age (since application) for different grant lags, computed for 12 years after grant

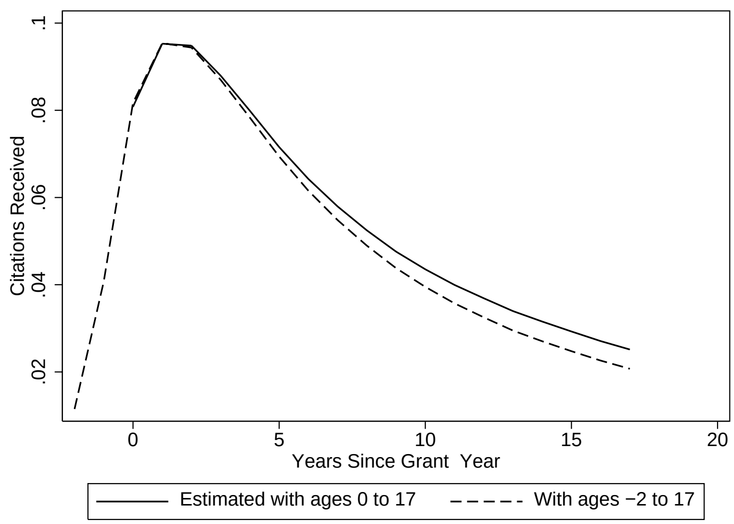


Figure 5: Age profile of citations

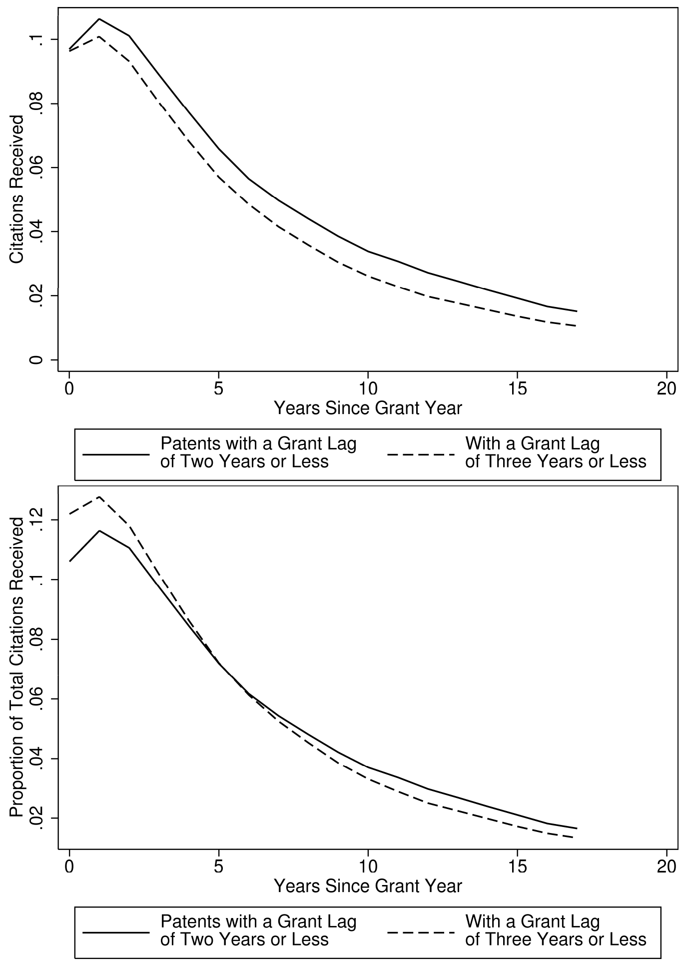


Figure 6: Changes in the age profile when adding patents with longer lags

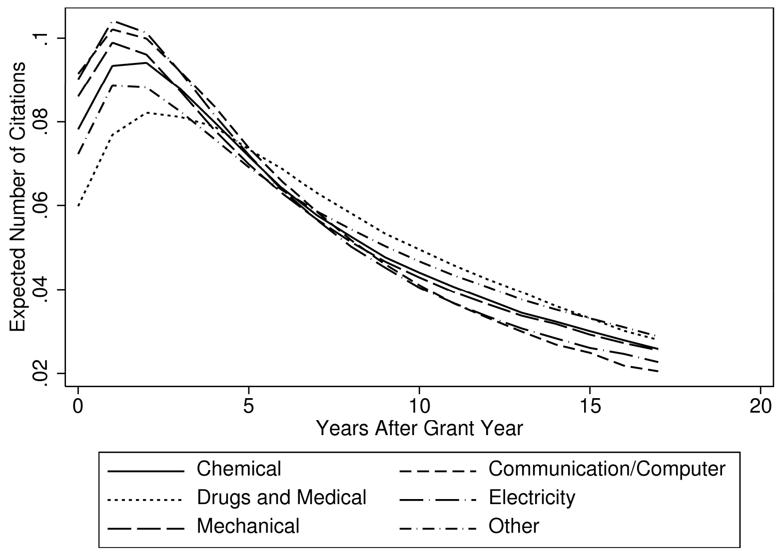


Figure 7: Difference in age profiles across industries

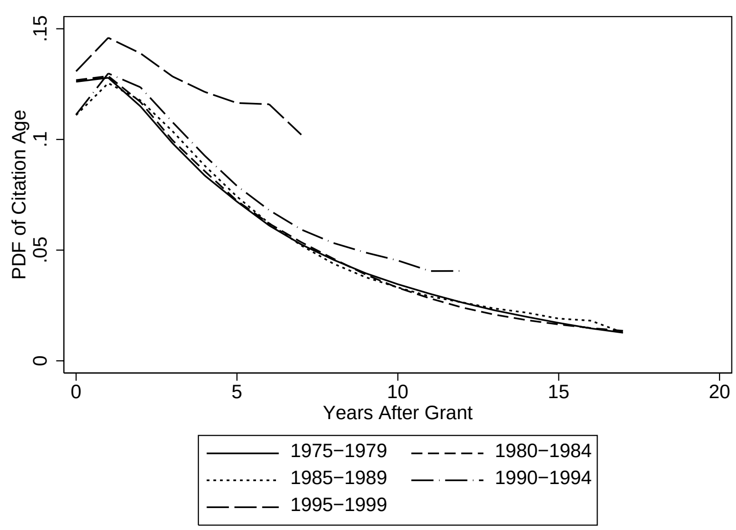


Figure 8: Difference in age profiles across grant-year cohorts