The Impact of Innovative Public Procurement on Technological Generality: a Patent Data Analysis

Emilio Raiteri[a]

[a]*Department of Economics and Statistics "Cognetti de Martiis", University of Turin*

Abstract

Keywords: *General Purpose Technologies, Public Procurement, Patent Data, Conditional Difference-in-differerence*

1. Introduction

Scholars have long acknowledged the important role played by market demand in shaping technological change and setting the pace of innovation (Schmookler, 1962; Kaldor, 1966). Though the difficulties in clearly disentangling between supply-side and demand-side induced innovations (Mowery and Rosenberg, 1979; Dosi, 1982) slowed down the study of this relation, the demand-pull hypothesis was never abandoned and is recently regaining momentum. In this context also the debate on the influence of public demand on technological change has received growing attention. In particular, both economists and policy makers are increasingly considering the innovative public procurement as an effective form of public support to private innovation activities, grounding the need for demand oriented technology policy (Edquist and Hommen, 2000b; Edler and Georghiou, 2007).

While the acknowledgment of public procurement as a *de facto* technology policy by policy-makers is a recent story, economic historians have long suggested an even more fundamental role for public procurement in setting the pace of technological change. Several works that studied the technological evolution in the United States in the 20$^{\text{th}}$ century, and mainly after World War II, stressed how the government demand has been a crucial factor in developing the most influential technologies of the last sixty years (Mowery and Rosenberg, 1982; Levin, 1982; Langlois and Steinmueller, 1999). In particular, Ruttan (2006) reports how U.S. military and aerospace related procurement had a major impact for the emergence and diffusion of every general purpose technologies (GPT) developed in the U.S. during the last century, such as semiconductors, computing and the internet. Even though it is common opinion that the end of the cold war reduced the chance for public demand to foster major technological breakthrough, recently numerous works (Antonelli, 2010; Mowery et al., 2010; Mazzucato, 2013) put forward that state intervention, and public procurement in particular, may have a very important role in spurring the green technology revolution and addressing grand social challenges.

Also the theoretical literature dedicated to GPTs, and in particular the seminal work of Bresnahan and Trajtenberg (1995) (BT), suggest that public procurement may have an important role in affecting the arrival of a new GPT. According to BT, while GPTs could be thought as engines of economic growth, this potential growth is achieved only if the virtuous cycle of innovation complementarities is triggered between the sector that unveils the new technology (upstream sector) and the sectors applying the new technology (downstream). Technological levels of the upstream and downstream sectors are hence strategic complements and widespread diffusion stems from the coordination of beliefs between the GPT producer and the application sectors. Coordination failures and larger uncertainty tied to drastic innovations may therefore provide little market incentives for adoption in the downstream sectors, potentially leaving an economy locked-in on inferior technological trajectories. BT hence already suggested that public procurement may play a very important role to overcome market failures, injecting the virtuous cycle through the stimulus of additional innovation complementarities.

Despite the economic historians' contributions and Bresnahan and Trajtenberg (1995)'s suggestion, no empirical work has so far tried to find econometric evidence of the link between public procurement and technological generality. This paper tries to fill this gap.
Following the intuition provided by BT and conceiving the arrival of a GPT "as a process unfolding in time rather than a single homogeneous shock" (Cantner and Vannuccini, 2012), I surmise that procurement might represent one of the most important element in creating the right soil to "cultivate" a technology that may (or may not) have the potential to reach high levels of pervasiveness.
To formalize this hypothesis, I make use of patent data and in particular of patent citations. Citations allow me to identify the connection between innovations related to public procurement and their patented antecedents, and to measure the degree of pervasiveness of a patent looking at the extent to which the follow-up technical advances are spread across different technological fields, through a Generality Index (Trajtenberg et al., 1997). On the basis of these considerations I will therefore hypothesize that receiving a citation from a patent related to public procurement raises the generality level of the cited patent with respect to the counterfactual situation in which that specific citation does not arrive.

In order to corroborate my hypothesis I hence design a quasi experiment in which we compare the change in the generality level (measured through the Generality Index) at two different points in time, 1999 and 2006, between treated and a control patents, whose application date falls in the period 1993-1997. Public procurement is the treatment variable and, in particular, a patent is put into the treatment group if it receives a citation from a patent related to public procurement in 1999-2000. To build the relevant variables for the quasi-experiment I create an original dataset exploiting data from four different sources: i) NBER patent data project that collects data for patents granted by the USPTO in the period 1976-2006, together with citations data; ii) Federal Procurement Data System (FPDS), that includes several information for each Federal contract awarded from 2000 onwards; iii) USPTO patent full-text and image database, which offers the full searchable text of every patent granted from 1976 onwards; iv) the Compustat North America Database that gathers financial and market information on public companies in the United States .

Since I clearly do not observe treated patents in the counterfactual situation in which they do not receive the procurement related citation, I have to use non-treated patents to proxy for it. If I was to use the whole non-treated paper as counterfactual patents, taking simple difference in averages of the change in the Generality Index between the treated and the non-treated patents would lead to biased results due to multiple endogeneity issues and mainly due to selection bias. To mitigate this problem I hence use as control patents only patents that are similar to the ones in the treated group along several dimensions[1]. I therefore adopt the conditional difference-in-differences (CDiD) approach developed by Heckman et al. (1998).

The estimate of the average treatment effect retrieved through the CDiD approach suggests a positive and significant impact of innovative public procurement upon the generality of a patent. In particular, on average receiving a citation by a patent related to public procurement raises the Generality Index of a 3.6%, confirming the initial hypothesis.

In the next section I briefly describe the different strands of literature that provide the motivation and the rational for this work. In section 3 I present the formal hypothesis tested in this paper, while in section 4 I illustrate the data and methodology I use in the empirical analysis. Section 5 presents results and several robustness checks. Conclusions follow.

2. Theoretical framework

2.1. Innovative public procurement as a technology policy

The idea that demand might be a major source of technological change dates back to the seminal contribution of Schmookler (1962) and Kaldor (1966). Despite the slowdown in the study of this relation occurred in the 80's due to the disruptive critics by Mowery and Rosenberg (1979) and Dosi (1982), the demand side approach has slowly but constantly regained attention (among others: (Von Hippel, 1988; Malerba et al.,

[1] As it was recently done in Czarnitzki et al. (2011), Feldman and Yoon (2012), and Fier and Pyka (2012)

2007; Rogers, 1995; Fontana and Guerzoni, 2008)). With the resurrection of the demand side also the debate on the role of public demand in fostering innovation has been revitalized.

Even if the impact of governments demand on firms' behaviour may appear evident just considering its size[2], particular attention has been recently given to its technological and innovative composition and to what is usually classified as 'innovative public procurement' [3].

Innovative public procurement is generally considered to occur when 'a public agency places an order for a product or a system which does not exist at the time, but which could probably be developed within a reasonable period'. This form of purchasing is usually opposed to 'regular public procurement' which occurs when a public agency buys ready made simple products such as pen and papers, where no R&D is involved (Edquist and Hommen, 2000b). However, recent works (Uyarra and Flanagan, 2010; Rolfstam, 2012) highlighted the potential limitations of this simple definition and stressed the fact that constraining the scope of innovative procurement to what happen after the placement of a formal order from a public agency is missing potential indirect effects of procurement on firms' behavior. In this paper, as in Guerzoni and Raiteri (2013), I hence follow a somewhat broader definition that deem as innovative public procurement all the 'purchasing activities carried out by public agencies that lead to innovation" (Rolfstam, 2012).

Leaving aside the debate on a narrower or broader definition, which is not the main concern of this work, over the last years innovative public procurement has been increasingly considered as a form of public support to private innovation activities, and hence as a 'de facto' technology policy (Cozzi and Impullitti, 2010)[4]. Several theoretical works (Geroski, 1990; Dalpé, 1994; Edquist and Hommen, 2000a; Edler and Georghiou, 2007) emphasized the potential positive effects of innovative procurement upon firms' innovative behavior through multiple and interacting channels. In the first place, public procurement is in fact thought to provide a minimal market size that allows firms to compensate costs and reduce the risks involved with R&D investments on products or services for which private demand is highly unpredictable. As we will discuss in section 2.2, this effect may be mostly important in the case of radical innovations whose development is usually characterized by larger uncertainties and arduous risk evaluations (Helpman and Trajtenberg, 1994). Secondly, public agencies may act as lead users providing producers with precious information about market needs and requirements, and also enabling firms to uncover already existent demand unmet by current products or services. Moreover, public procurement can represent a very useful tool in standard setting and diffusion of specific technologies.On this ground numerous scholar called for the need of 'de jure' procurement oriented innovation policy. A call that is recently receiving more and more positive answers at the political level as it is well testified by several documents issued by European Commission (EU, 2010) and the OECD (OECD, 2013), in which innovative public procurement is acknowledged among other more consolidated technology policies, such as R&D subsidies and tax credits.

Along with theoretical and political attention, a growing body of literature providing quantitative empirical evidence about the positive impact of public procurement on firms' innovative behaviour is joining the abundant qualitative evidence reported in case studies(Edquist and Hommen, 2000a; Rolfstam, 2009; Uyarra and Flanagan, 2010; Flanagan et al., 2011; Brammer and Walker, 2011). An early work in this area by Lichtenberg (1988) tested the effect of federal procurement upon contractors' private R&D expenditures. His result suggests that public procurement not only has a positive effect on a firm's propensity to engage in R&D, but also that the demand pull effect is larger for public procurement than private contracts. A more recent paper by Aschhoff and Sofka (2009) test the role of various technology policies (R&D subsidies, innovative public procurement, regulation, university research) on a cross-section of 1149 German firms[5]. They compare the impact of each policy on firms' innovative output, proxied by the share of turnover with market novelties. They find robust evidence for a positive impact of public procurement, in particular for small size firms. Guerzoni and Raiteri (2013), using data on 5238 European firms from the 'Innobarometer on

[2]According to OECD (2013) member countries spend on average 13% of their GDP on public procurement

[3]Expressions like 'public technology procurement' and 'public procurement of innovation' are used to refer to very similar phenomena. For further discussion see Rolfstam (2012)

[4]For a review of the state of the art of this debate on innovative public procurement, including definitions and taxonomies, see Uyarra (2013)

[5]Firms that responded to the survey 'Mannheim Innovation Panel' in 2003

Strategic Trends in Innovation 2006-2008' survey, provide new evidence about the impact of three different technology policies, innovative public procurement, R&D subsidies, and tax credits upon firms' innovative behavior measured in terms of innovative input (R&D expenditures). Their results suggest that innovative public procurement is a very effective tool in raising private expense in R&D, especially when administered together with other complementary technology policies.

2.2. Public procurement in the economic-historical analysis of technological change

As pointed out in the previous section, public procurement is nowadays increasingly considered as an effective policy tool for fostering innovation by policy-makers and innovation scholars. However, most of the recent works on innovative public procurement say little about the kind of innovations that public procurement is able to induce and about their technological impact. This gap is most striking if we consider the consistent amount of historical and economic analysis devoted to investigate the role played by defense related procurement in shaping the patterns of technological change during the 20th century, especially in the United States. Some of the most interesting works in this field are already collected in a volume edited by Richard Nelson in 1982 (Nelson, 1982), in which different scholars analyze how public policies affected technical progress in seven key American industries[6]. The contributions of Levin (1982), Mowery and Rosenberg (1982), and Katz and Phillips (1982) stress in fact how the sheer size of procurement for components and systems for purposes of national defense and spatial exploration drew forth fundamental technological advances in the semiconductor, the computer, and the aviation industry.

Levin (1982) highlights that the presence of government demand abundantly reduced the risk of investment in semiconductors technologies such as the silicon transistors and the integrated circuit in the early years of their development. Several further works (Mowery and Rosenberg, 1989; Langlois and Steinmueller, 1999; Mowery, 2011, 2012) confirmed that vast procurement contracts drove R&D private effort in the semiconductor sector and also that some of the most important breakthroughs in the industry, though achieved by privates, were undertaken with the needs of the military foremost in the minds of the successful inventors. According to Levin (1982) in 1959-60 federal procurement absorbed between 45 and 50 per cent of the total semiconductors industry output and more than 50 per cent of the productions of integrated circuits from 1962 to 1966. Indeed, the prospect of large procurement contracts appears to have operated as a prize leading potential contractors to invest their private funds to develop new products that met government demand requirements (Mowery, 2011).

In the same way, also in the computer industry, the federal procurement accounted for more than 50 percent of total shipments between 1945 and 1955 (Flamm, 1987), and, even though the federal share declined substantially in the late 1950's for the emergence of private demand, governmental demand still represented more than 40 percent of total sales of supercomputers at the beginning of the 1970's. As in the case of semiconductors, contracts between government and private firms therefore had a profound influence in shaping the structure of the nascent computer industry between 1945 and 1960 (Katz and Phillips, 1982) and the sheer size of defense related procurement seems to have acted as a powerful attractor for new firms to enter the industry and develop new products and applications Mowery (2011). (Mowery and Rosenberg, 1982) see a similar pattern in the rise of the U.S. aircraft industry, for which the existence of government demand was crucial in bringing about rapid diffusion of new technological knowledge.

While the afore mentioned studies consider, in most of the cases, public demand as one of the multiple factors that facilitated diffusion and improvements in specific key technologies, Ruttan (2006) pushes the argument even further. He claims in fact that defense related procurement has been the most important factor for the development of every general-purpose technology in which the United States was internationally competitive throughout the 20$^{\text{th}}$ century, from the deployment of interchangeable parts and mass production system, passing through commercial aircraft, nuclear energy, semiconductors and computers, to arrive to the building of the internet, space communication and earth observing technologies. In particular Ruttan speculates about the potential counterfactual situation in which the military demand was not there to actively create markets for those technologies, and he advocates that, in each of these cases, commercial development would have been at least substantially delayed without the governmental stimulus.

[6]Semiconductors, Commercial aircraft, Computers, Agriculture, Pharmaceuticals, Motor Vehicles, Residential Construction

From the beginning of the'90s several studies argued that changes in the global political situation due to the end of the "cold war", together with the changes in the structure of the U.S. economy, precluded defense and space related procurement from continuing to play an important role in developing radical technologies that might have "dual use", both military and commercial. Even though this might be true for matured technologies such as the semiconductors or the computer, Cowan and Foray (1995) stress how this view suffered from over-generalization since it did not consider the lifecycle of a technology. According to Cowan and Foray (1995), defense-related R&D can be able to play an important role for the development of emerging technologies also in the post-cold war era. They indeed put forward that, when a new technology arrives, the scope of its technological impact is hard to predict, and that both military and civilian sectors exhibit a similar degree of ignorance about its future trajectory. Civilian R&D, focusing on profitable applications only, will explore a different portion of the technological spectrum with respect to military R&D, which will be instead more interested in technical dimensions and performance rather than costs. Defense-related R&D therefore increases the diversity of applications of a new technology and, even more important, increases the diversity of information available about the emerging technology (Cowan and Foray, 1995). More information about a new technology reduces both the costs and the ambiguity attached to further innovations, fostering development and diffusion also in the civil market .

Very recently some policy oriented works started again to consider the public agencies' purchasing activities as an important source of major technological breakthrough. Antonelli (2010) advocates that raising consistently the technological composition of public demand not only for defense-related scope but also for education, health, energy may trigger those beneficial interactions that lead to bandwagon of radical innovations and to the deployment of new technological systems. Mazzucato (2013) underlines that the state, being less risk-averse than private sector, has always played and still play a fundamental role in fostering radical growth-enhancing innovations. Mazzucato also suggests that public intervention might be particularly effective and desirable in the present-day in order to achieve the green-technology revolution, a process requiring huge investments that we can not expect from the private sector alone. On the same ground Mowery et al. (2010) and Foray et al. (2012) put forward that public procurement might be effectively used to encourage the development of climate-friendly technologies.

2.3. General purpose technologies and innovative public procurement

The historical evidence depicted in this section seems then to support the existence of a strong relationship between public procurement and the impact, in terms of adoption and pervasiveness, of fundamental technologies. More recent literature also suggest that public demand may still play a key role a key role in reaching new technological revolutions and the deployment of new technological systems. To better understand the nature of the link between procurement and the emergence of radical innovations and to comprehend its economic rationals, it is now useful to consider carefully the branch of literature that theorized the concept of general purpose technology, explicitly mentioned by Ruttan (2006) in its analysis.

The notion of general purpose technology was first introduced by David (1989) and thoroughly developed by Bresnahan and Trajtenberg (1995)(BT). Since then it has been used to refer to specific key technologies that shaped the process of technical change and productivity growth in different eras, such as the steam engine, the factory system, the electricity, and semiconductors[7]. Bresnahan and Trajtenberg (1995) and Rosenberg and Trajtenberg (2004) formally define a GPT as a technology having the following characteristics: i) General applicability, or General purposeness, which means the technology should perform some generic function that is vital to the functioning of many products or production systems in downstream sectors; ii)Potential for continuous technical advance in the efficiency of the generic function after its introduction; iii) GPT exhibits "innovational complementarities" with the application sectors, which means that advances in the technology foster innovation in the downstream sectors, magnifying the technological impact of an innovation in the upstream sector.

[7]In this respect the concept of general purpose technology is not to distant from other ideas that tried to take into account the uneven nature of technological change such as radical innovations (Schumpeter, 1934), technological paradigms (Dosi, 1982), macro-innovations (Mokyr, 1990). For a more in depth analysis of the relation among these different notion see Cantner and Vannuccini (2012).

GPTs can therefore be thought 'engines' of economic growth, able to transform the inner structure of an economy, achieving raise in productivity and output in the long run. Even though the literature that followed BT focused on this latter feature of the GPT and used the concept mainly to construct endogenous-growth models (Cantner and Vannuccini, 2012), BT underline that this growth enhancing transformation is achievable if and only if a dual inducement mechanism between the GPT sector(upstream) and the application sectors(AS) (downstream) is set in motion. This process is mostly related to the third of the GPT's features listed above, and is in fact named 'innovation complementarities virtuous cycle'. It entails that improvements in the quality of the GPT foster R&D investments and innovations in the application sectors, which, by rising the technological level in the application sector bring on further investments and technical improvements in the GPT sector, from which, in turn, stems further adoption in the AS, and so on and so forth. The quality of the GPT and the technological level in the AS can hence be seen as strategic complements in the process of generating innovation complementarities and in determining the size of R&D investments in the upstream and downstream sector. The sketch in figure 1 tries to summarize this intuition of BT.

[Figure 1 about here.]

The arrival of a GPT and its widespread adoption in different sectors derives from the coordination of beliefs about the potential development of the technological trajectory between the GPT sector and the application sectors. BT show in a game-theoretic model that the successful coordination among sectors is not granted in a decentralized market system due to the existence of two different kind of externalities. The first one is a vertical externality that link the payoff of the firms in the application sectors and in the GPT sector. Since both side would like to appropriate the social returns coming from the deployment of the GPT, neither the upstream nor the downstream sector will have the incentive to innovate. The second externality is horizontal and takes place across application sectors. Since the more AS adopt the GPT the higher will be its quality and, therefore, also the larger incentive to adopt it in the AS. The technological level in the GPT sector, being a function of the adoption in the AS, acts as a 'public good' but no firm in the AS will have the incentive to contribute to its production, since only the other firms would reap the benefit of that effort. Moreover, the uncertainty attached to technological change, and in particular to radical innovations (Rosenberg, 1998), exacerbates the coordination problem that the externalities bring about. BT model puts forward that uncertainties about technological trajectories and imperfect appropriabilities may lead to coordination failures in determining the optimal level of R&D investment both in upstream and downstream sectors, hindering or delaying consistently the realization of the innovation complementarities virtuous cycle that enable the potential arrival of a GPT. BT also suggest that this sort of coordination failure call for policy consideration and for government intervention to fix the market failures. In particular, drawing from historical evidence and looking at the role played by the U.S. Department of Defense and NASA during the fifties and the sixties, they explicitly suggest that public procurement may set in motion and sustain the innovation complementarities virtuous cycle. Public procurement may in fact create and enlarge markets and therefore stimulate private investment in R&D, innovation and adoption in the application sectors on a scale that would not have otherwise followed.

[Figure 2 about here.]

As the sketch in figure 2 once again tries to summarize, the procurement induced innovation complementarities may in turn stimulate investments in R&D in the upstream sector, triggering the virtuous cycle that will potentially lead to the unfolding of a new GPT.

3. Research hypothesis and identification strategy

While the historical literature provided some qualitative evidence of the existence of a direct relations between innovative public procurement and the degree of pervasiveness of specific technologies, the work of BT on GPT offered instead some theoretical ground for understanding it. As discussed in the previous section, even though the GPT literature then followed another path, BT did not see the deployment of a GPT in terms of technology arrival but rather as a process unfolding in time. Recently Cantner and Vannuccini (2012), reconsidering closely the most important works on GPTs, put forward that the degree of generality of a technology should not be considered as an ex-ante characteristic, steady over time, but as dynamic attribute that may evolve. In particular they suggest that GPTs can be, at least to some extent, "cultivated or developed [..] till they assume the role of core technologies". Not only *A time to sow and a time to reap* then, as described by Helpman and Trajtenberg (1994), but also a time to *nourish* technologies in the early phase of their development. On the ground of these contributions, the main idea of this paper is that, in this 'technology cultivation process', innovative public procurement can be a key sustenance to increase its pervasiveness and, therefore, its level of generality.

There are several reason that suggest the relevance of public procurement in this respect. As already mentioned, for a technology to become very general it takes that different application sectors adopt it and develop new products embodying it, inducing further improvements in the upstream sector that in turn will spawn further adoption. Studies on adoption and diffusion of new technologies (Geroski, 2000; Nelson et al., 2004; Hall, 2006) suggest that four main factors mostly affect users' adoption decision: (i) performance of the new technology; (ii) costs of adopting the new technology; (iii) network effects; (iv) uncertainty. According to Carlaw and Lipsey (2011) the latter element could be particularly important for the diffusion of radical technologies such as GPTs, since both the application and the GPT sector would act in condition of knightian uncertainty (or ambiguity) rather than risk. An evaluation of the current and future performance of the new technology and of its adoption costs will heavily depend on learning processes and on the amount of knowledge available about the new technology, which in turn will depend on the coordination of beliefs between potential adopters and the upstream sectors. In the early phase of development of a new technology, application sectors will not be able to compute probabilities for the potential outcomes since the latter will rest on the results of highly uncertain process such as R&D investment and knowledge production, learning, and coordination. In the absence of spontaneous coordination, private firms may therefore explore only a limited portion of the variety distribution of a technology (Foray, 1997), or even completely avoid to invest in the new technology and keep doing R&D for improving on existing technologies for which is possible to compute expected returns on the investment. In such a context, public procurement is hence a powerful tool to foster adoption and innovation since, in the first place, it is able to absorb most of the uncertainty related to profitability and costs that the private sector is not willing to face (Mazzucato, 2013). The primary interest of the state in its procurer activity is not profit but the satisfaction of specific needs that could transcend the short term economic feasibility of a project. One obvious example is national security in the U.S. during World War II, or in the cold war era. The state, through public procurement, can hence create a market large and profitable enough to stimulate abundant investment in R&D and to develop innovation complementarities based on technologies whose success is too unpredictable for the private sector alone, or that explore a different portion of the technology spectrum distribution (Cowan and Foray, 1995; Fabrizio and Mowery, 2007; Mazzucato, 2013). Moreover, the state, through procurement contracts, can directly take charge of the cost-related uncertainties. For instance, in the United States the Federal Acquisition Rules suggest that, when "uncertainties involved in contract performance do not permit costs to be estimated with sufficient accuracy", cost-plus or cost-reimbursement contracts should be considered as suitable contracts. With such kind of contracts, a public agency in order to have a specific product or service delivered, reimburses the contractor the realized (or a share of the) cost and pays an additional fee[8]. Given the inherent uncertainty that characterize the research activities, cost-plus contracts are used for most of the federal procurement contracts that involve the performance of R&D (from basic research to development). Public procurement is therefore able to foster innovations that would be otherwise inhibited by the high degree of ambiguity that permeates the evolution of new technologies in the early phase.

These procurement-related innovations contributes to set the complementarities virtuous cycle in motion in two ways. On the one side, innovations that apply a new technology generate new knowledge and new information about the new technology itself. In an environment characterized by knightian uncertainty,

[8]The additional fee might be fixed ex-ante or based on performance

enlarging the amount and the diversity of available knowledge reduces the uncertainty associated with further innovation, focusing subjective probabilities on those inventions that are more likely to succeed (Bewley, 2001). The reduction of uncertainty then act as positive feedback, stimulating other innovations as new information becomes available. Innovative public procurement hence may convert knightian uncertainty into risk, allowing more ambiguity-averse firms to reasses their expected returns on innovation and on the adoption of the new technology (Geroski, 2000). On the other side, procurement-induced innovations act as a coordination device between the upstream and the downstream sectors. By creating or enlarging markets in the application sector it will in fact also stimulates further investment in R&D in the upstream sectors that will improve the performance-price ratio in the upstream technology, favoring further diffusion and innovation in the downstream sectors and setting in motion the virtuous cycle that can lead to the arrival of very pervasive technology.

Clearly I will not argue here that public procurement support to a specific technology can always lead to very general technologies. Finding evidence of the arrival of new GPTs is therefore beyond the scope of this paper [9]. I will instead put forward that public procurement, by stimulating additional innovation complementarities in the application sectors, will increase the probability of diffusion of upstream technologies among different sectors, making them more pervasive, or general, compared to the counter factual situation in which no stimulus from public procurement was in place.

I will neither hypothesize that public procurement might be able to directly produce major innovations, as it has been suggested in some cases (Ruttan, 2006), but that it can support promising new technologies early in their life-cycle (Levin, 1982) through the stimulus of innovative activity in the AS, that would not have occurred in the absence of public demand.

Despite the numerous contributions by economic historians and the theoretical work in the GPT literature, no econometric work tried to investigate the effect of innovative public procurement on the technological impact in term of pervasiveness of given technologies. This paper tries to fill this gap through a patent data analysis and exploiting public data on Federal Procurement made available by the U.S. Government. In order to formalize my research hypothesis, in the next section I will explain why and how patent data can be use for identification.

3.1. Patent data and patent citations

Patents are transitory monopolies granted to inventors/assignees for the commercial use of a new product or process in exchange of full disclosure; they are hence usually considered as direct output of the inventive process and, more specifically, they represent outcomes that are expected to have an economic and commercial impact. Moreover, being legal documents, they have to include several detailed information on the innovation, the inventors, the assignees and prior art. For these reasons, patent data have long been acknowledged as a very important source of information for researcher studying technological change (Griliches, 1991).

The digitalization of patent documents in the 1980's and the continuous improvements in computational power throughout the 1990's helped scholars to extract more information from patent data, and to make it even more valuable for empirical research. In this context, several works highlighted the fundamental importance of the information embodied in patent citations for determining the economic value and the technological impact of a patent (Trajtenberg, 1990; Jaffe et al., 1993; Trajtenberg et al., 1997; Jaffe et al., 2000; Hall et al., 2006).

The most important feature of patent citations is the fact that they have a fundamental legal value. An innovation to be patentable has to be novel, non-obvious and useful. The degree of novelty and the scope of the property right awarded through a patent are hence delineated by the citations to the technological antecedents of the invention because, if a patent cites another one, it means that the latter constitutes a piece of previously existing knowledge upon which the former builds and over which it cannot have any claim

[9] Several papers tried to detect the establishment or the emergence of new GPTs. See for example Hall and Trajtenberg (2004), Youtie et al. (2008), Feldman and Yoon (2012).

(Hall et al., 2006). Citations are hence mandatory and identify the relevant prior art[10].
Given their legal function, citations actually represent the stream of the past relevant knowledge that feeds the production of the new pieces of knowledge. It is hence possible to consider citations as a "paper trail" (Hall and Trajtenberg, 2004) of the linkages between an innovation and its technological antecedents (citations made/backward citations) and descendants (citations received/forward citations). Because of these features scholars consider citations as carrier of different kind of informations about a patent, such as its economic value and quality (Trajtenberg, 1990; Hall et al., 2006), the knowledge spillovers (Jaffe et al., 1993) it involved and its technological impact(Trajtenberg et al., 1997; Henderson et al., 1998).

In the context of this paper we are primarily interested in the latter aspect. Trajtenberg et al. (1997) created different measures based on patent citations that tried to capture both the basicness and the technological importance of an innovation covered by a patent. The main idea is to exploit the fact that patent offices, in order to facilitate the quest for prior art, assigns patents to specific technology class and subclass (about 400 (3-digit) classes and 120,000 subclasses). Citations between different inventions could be hence considered also as linkages between different patent classes and hence technological fields. Trajtenberg et al. (1997) suggested that a patent that receives forward citations from patents belonging to wide range of different classes could then be thought as a patent having a wide variety of application, and therefore as a very generic patent. To operationalize this intuition they developed the, so called, Generality measure, an index that get closer to 1 as a patent's forward citations are spread across many different patent classes, and approaches 0 as forward citations are concentrated in a few technological classes. Though this index will be presented in more details in the next section, it is important to note here that it allows us to measure effectively the degree of pervasiveness of given innovations covered by patents and to compare it across patents and time. While Trajtenberg et al. (1997) used this measure to evaluate its impact on the degree of appropriability of an invention, several works (Hall and Trajtenberg, 2004; Moser and Nicholas, 2004; Youtie et al., 2008) used it effectively to uncover the existence of general purpose technologies or to spot the deployment of specific technological trajectories, confirming that patent and patent citations can be particularly fit to assess the degree of generality of a technology.
Moreover, a more generic feature of citations is pivotal for my empirical analysis. As mentioned above, citations provide the link between present inventions and previous inventions (Hall and Trajtenberg, 2004).Exploiting this feature I can therefore use patent citations to spot the bond between innovations stimulated by (related to) public procurement and their patented technological antecedents. I then interpret the latter as upstream technologies and the former as innovation in the application sectors. Therefore, patent and patent citations, through the detection of the link between innovations together with the generality measure, allow me to put the broad idea portrayed in the previous section in a more formal way. I will in fact hypothesize that:

> *Hyp: Receiving a citation from a patent related to innovative public procurement (i.e. an additional innovation complementarity covered by a patent) will rise the degree of generality of the cited patent (upstream technology) compared to the counterfactual situation in which that specific citation does not arrive*

Patents will hence be the unit of analysis throughout the paper, and the focus variable will be (the change in) their generality level measured through patent citations. In order to test my hypothesis, I will frame the problem as a quasi-experiment in which the generality of a selected group of patents will be measured at two points in time. I will then compare the change in the generality level across time between a group of treated and group of control patent. A patent will be assigned to the treated group if it receives a forward citations from a patent related to public procurement right after the first measurement of its generality level. The control group will be instead carefully constructed to approximate the counter factual situation in which the treated patents do not receive the treatment, as it was done in several recent works (Lanjouw

[10]For applications at the U.S. the applicant has in fact the duty to disclose any knowledge of the prior art of which he is aware, and then the patent examiner, an expert in the area, certifies that all the relevant prior art have been included and eventually adds missing citations. The same rational holds at the European Patent Office (EPO) but all citations are added by the patent examiners who add the minimum number of citations to cover prior art (therefore patents granted in the U.S. usually report more citations than the EPO ones due to the different institutional procedure).

and Schankerman, 2004; Czarnitzki et al., 2011; Feldman and Yoon, 2012; Fier and Pyka, 2012). Figure 5 depicts the basic idea of the natural experiment proposed here.

[Figure 3 about here.]

4. Data and Method

4.1. Data:

To build up the database that I use in my analysis I put together information coming from four different sources. The NBER patent database 2006 is the main source of information. Providing data on patents and citations, it allows me to identify a large sample and to construct the outcome variable, i.e. generality index. In the second place I exploit the public data from the U.S. Federal Procurement Data System together with USPTO database, to single out patents related to public procurement, and hence to construct my treatment variable. Finally I also employ the Compustat North America database in order to gain more information on the patent assignees and to better select patents to be included in the control group. In this section I briefly describe the broad information available in those databases, while in section 4.2 will explain how I used them to build up the variables and the framework of my natural experiment.

4.1.1. NBER Patent Database - 2006

The NBER patent database[11] contains information on 3,209,376 unique patent granted by the USPTO from 1976 to 2006[12]. Hall et al. (2001), who developed the dataset, carefully describe the information included in its first version, which has been continuously updated since. Several informations are available for each patent: the patent number, the year in which the inventor applied for the patent (Application year), the year the USPTO granted the patent (Grant year), the country (and the state if U.S.) of the inventor, the assignee identifier and the type of assignee (individuals, U.S. corporation, foreign corporation, governments, university), the main U.S. 3-digit patent class (400), the subclass (120,000) and the number of claims made by each patent. Moreover additional informations are made available in different complementary files that reports data on the inventors, the full name of the assignee, an identifier that allows the matching with the Compustat North America Database and the citation data. The latter file include 23,650,891 references between cited and citing patents (listed through their unique patent number), and the total number of citations received by the cited patent at the end of the screened period, in 2006.

4.1.2. USASPENDING.GOV and Federal Procurement Data System

In 2006 the U.S. congress approved the Federal Funding Accountability and Transparency Act (FFATA), sponsored by Senators Coburn, Obama, Carper, and McCain. The Act required the Office of Management and Budget (OMB) to establish a single searchable website, accessible to the public at no cost, which includes for each Federal award: the name of the entity receiving the award; the amount of the award; information on the award including transaction type, funding agency, etc; the location of the entity receiving the award; and a unique identifier of the entity receiving the award. In order to fulfill these requirements, in December 2007 the U.S. government launched USAspending.gov [13], a website that collects prime award data for federal contracts, grants, direct payments and loans. The most interesting feature of this dataset, at least from the point of view of this paper, is that it includes full Federal Procurement Data System (FPDS) from Fiscal Year 2000 (October 1999) onwards. The FPDS tracks every public procurement contract over 3000 dollars between federal agencies and contractors. Several pieces of information are available for each contract and in particular: the obligated amount of the contract, the purchasing agency, the contractor, a code describing the product or service being purchased, the kind of contract (cost-plus or fixed cost), the extent to which

[11] Data and data description are available at `www.sites.google.com/site/patentdataproject`.

[12] The NBER patent dabase is the result of the effort of different researchers, Hall et al. (2001) beside the authors, credits Rebecca Henderson and Michael Fogarty, together with several programmers and research assistants.

[13] More information and data are available at `www.usaspending.gov`.

the contract was competed. To have a clearer idea of the size and characteristics of federal procurement in the U.S. we can have a quick look at some aggregate data for Fiscal Year 2000, the period that I will refer to in this analysis. In FY 2000, government agencies awarded 594,541 contracts for more than 205 billion dollars. Almost 65 per cent of the total amount allocated was awarded by the Department of Defense(DoD), 10 per cent by the Department of Agency(DoE), 5 per cent by General Service Administration (GSA), 2.5 by the National Aeronautics and Space Admnistrations (NASA), and then less than 2 per cent each by 54 other governmental agencies. About 6 per cent of the total contracts were awarded for the performance of some kind of R&D (from basic research to development). 78 per cent of the contracts was assigned through competition, while 22 per cent did not involve any competitive bidding activity. These numbers are in line with the ones for the whole period 2004-2010 described in Liebman and Mahoney (2013). FPDS data report that the total number of federal contractors who won at least one procurement contract in fiscal year 2000 is 47,084, nevertheless the allocation of resources between them is highly skewed. The top 10 contractors[14] account for slightly more than 30 per cent of the total amount awarded through procurement contracts, the 50 largest contractors account for 48 per cent, and top 1 per cent (i.e. top 470 contractors) accounts for 70 per cent of the total procurement expense.

4.1.3. USPTO Full-text and Image Database

The USPTO Full-text and Image Database includes information about US patents from 1790 to the present day. The feature making this database extremely helpful lays in the fact it offers the full searchable text of every patent applied granted from 1976 onwards. In particular it possible to search for specific pieces of text within many distinct field of the patent document, such as the patent's title, the assignee's name, the abstract, the claims, the description of the invention, ecc.[15]. As we will see in the next section, the most relevant field for the purposes of this paper is the Government Interest field, which contains data describing the US Government's Interest and rights in the patent, if any.

4.1.4. Compustat North America Database

Compustat North America provides annual and quarterly financial and market information for more than 90,000 (both active and inactive) publicly traded firms in the United States and Canada from 1962 to present day[16]. It includes abundant information about income statement, balance sheet, statement of cash flows, and supplemental data items. In this work I will use only a limited number of variables from Compustat Database, and in particular those concerning the number of employees, the size of sales and net income, and the amount of investment in R&D of the firm.

4.2. Quasi experiment

I use the huge amount of information contained in the 4 dataset described in the previous section to design a quasi-experiment to validate my hypothesis. In the first place I hence identify the sample of patents that will be used in the analysis, and I will then carefully describe the building of the treatment and of the outcome variable.

4.2.1. Sample selection

From the NBER patent database I extract all patents whose application date belong to the five year window from 1993 to 1997. For each of these patents I aggregate the information about the backward citations they made, and the forward citations they received up to 2006, the last year for which data are available in the dataset. Exploiting the assignee and the compustat identifier provided by the NBER database (Hall et al., 2001) I match each patent with the Compustat North America Database in order

[14] Lockheed Martin Corporation, Boeing Company, Northropp Grumman Corporation, General Dynamics, Raytheon Co, State of California, Bechtel Group Corporation, BAE Systems, McDonnel Douglas, SAIC Company.

[15] A comprehensive list of fields is available at `www.patft.uspto.gov`.

[16] Specifically firms traded on the New York Stock Exchange (NYSE), American Stock Exchange (ASE), National Association of Securities Dealers Automated Quotations (NASDAQ), Over-the-Counter (OTC), Toronto Stock Exchange, Quebec Stock Exchange, and Montreal Stock Exchange. Though the project started in 1962, annual data are available back to 1950.

to gain additional information about the patents' assignees. This means that, as a first step, the scope of the analysis is limited to patents owned by public companies. In section 5.2 I will evaluate if the results are consistent when the breadth of the analysis is not constrained to patents owned by publicly traded companies. Moreover I only consider patent who received at least 10 patent citations at the end of the period, i.e. 2006. The rational behind this exclusion is twofold. In the first place since our outcome variable is constructed on the basis of the concentration of forward citations, looking at patents receiving only few references would be less interesting and could also introduce some bias: the fewer the citations, the larger would be the bias. As we will see, Hall (2005) showed how to correct for this bias and also that it quickly disappear as the number of citations increases. In the second place, as Scherer (1965) already hypothesized, the distribution of patents' economic value is highly skewed toward the low value side, with a very long tail into the high value side. The number of forward citations, which follows a Pareto-like distribution, proved to be a good predictor of a patent's economic and technological significance (Trajtenberg et al., 1997; Hall et al., 2006). In particular Hall et al. (2006) showed that patents receiving more than 7, and especially those receiving more than 10, forward citations are the ones carrying actual economic relevance for publicly traded companies. Therefore, focusing on patents with more than 10 forward references ensures that the results of the analysis will not be driven by patents with little economic and technological impact. We end up with a sample of 71,438 patents, for which we have detailed information about: patent number; year of application; grant year; the technology class (main U.S. 3-digit); number of claims; number of citations made; number of citations received; technology class of each citation made and received (3-digit); assignee typology (being public firms, basically if it is a U.S. or a non-U.S. corporation); sale, net income, industrial sector (SIC and NAICS code), and R&D expenses of the assignee. I also compute the originality level of each patent. Originality is a measure very similar to the generality index that will be carefully described in section 4.2.3. First proposed by Trajtenberg et al. (1997), it is meant to measure the level of originality of a patent on the basis of the backward citations he makes. In particular a patent wil be more original (Originality closer to 1), as it cites prior arts coming from many different patent classes (i.e. if it synthesizes divergent ideas (Trajtenberg et al., 1997)), while it will be less original if its backward citation are concentrated in a small number of classes.

To have a first look at the data, table 1 shows the distribution of patents in our sample across technological classes. As it is possible to see from the table, most of the patents belong to the ICT sector (40.7 per cent) and electronics (24.7 per cent), followed by mechanicals,chemicals and drugs.

[Table 1 about here.]

Table 2 reports instead descriptive statistics about each patent characteristics together with assignee specific characteristics.

[Table 2 about here.]

The patents in our sample spread over five years. The average patent in our sample makes 11.9 backward citations and receives 24.5 forward citations. The average lag between the application and the granting of the patent by the USPTO is about two years. As mentioned, forward citations are censored from below (10 cites), and, as figure 4 depicts, they are distributed in a very skewed way, with 56 per cent of the patents in the dataset (47,510) obtaining between 10 and 20 references, 22.7 per cent of them receiving more than 30 cites, 8.3 per cent more than 50, and only 1.3 per cent of the patents collecting more than 100. Moreover 3 out of four patents in our sample belong to U.S. corporation rather than foreign ones.

[Figure 4 about here.]

On average assignees are very large companies with more than 10,000 employees and whose sales reach 25 billion dollars. They also invest on average more than 4 per cent of their total revenues in R&D.

4.2.2. Treatment variable

In section 3 I hypothesized that public procurement, enlarging the diversity of application and available knowledge about a given technology through additional innovation complentarities, can increase its pervasiveness and its technological impact. In section 3.1 I propose to test this hypothesis through patent data, and, in particular, I surmise that a patent receiving a citation from a subsequent patent related to a public procurement contract will have a higher generality index compared to the counter factual situation in which no public procurement related reference arrives. Since the counterfactual situation is obviously not observable, I divide the sample in treated and control patents and I will use the latter to estimate a proxy for the counterfactual situation. As mentioned, a patent is put in the treatment group if it receives a citation form a patent related to public procurement, otherwise it ends up in the control group. Therefore, I now have to first define how I identify a patent related to public procurement, and, secondly, to build the treatment variable for my quasi-experiment exploiting this identification.

In order to identify the patent related to public procurement I follow a strategy in two step. In the first place I use the FDPS data for Fiscal Year 2000, a period that spans from October 1st 1999 to September 30th 2000. The reason to select only data for this specific year lies in the nature of available data for patent citations. Since the last year for which I have patent data available is 2006, and because I am mainly interested in the arrival of new citations after the implementation the treatment, I only use the first year for which I have procurement data accessible (FY 2000) to identify the treatment in order to maximize the period in which post-treatment citations may take place. As explained in section 4.1, FDPS data are recorded at the contract level and entails information about more than 590,000 contracts between federal agencies and private contractors. I aggregate the information coming from those contracts at the firm level and I then match manually through the entities name of each contractor to patent data in the NBER Patent Data Project. I hence select all the patent that have a priority date in year 1999 or year 2000 and belong to firms that won at least one public procurement contract in FY 2000. The rational to use priority date instead of application date in this context is that the former ensures to be closer to the actual date of invention and hence it allows to spot innovation closer in time to the award of the procurement contracts in FY 2000. Though the fact that a patent belongs to an entity that won at least one public procurement contract in FY 2000 and that it has been first filed to the USPTO at the same time is clearly not a sufficient condition to claim that a patent is linked to public procurement, it can be interpreted as a necessary condition. In this way I identify more than 32,323 patents potentially related to public procurement.

As a second step, to spot the patents that are actually linked to procurement among those in the group described above, I exploit the U.S. Federal Acquisition Regulation (FAR). The FAR is a set of rules that governs the purchasing of good and services carried out by U.S. federal agencies. Though FAR was first approved in 1974, it has been adapted to follow the Bayh-Dole Act of 1980 [17], for what concerns intellectual property right management in federal procurement contracts (Sharp, 2003; Bloch and Gray, 2012). One of the rational behind the act was to tackle the increasing reluctance of contractors to collaborate with the federal government due to "title taking policy" of many agencies (Sharp, 2003). In many cases, before the Act, the acquisition rules of federal agencies assigned the right to patent an invention realized by a firm in the performance of work under a government contract to the government itself, while the contractor could only obtain limited rights and licenses. The Bayh-Dole Act leveled the rules for the different agencies, granting more rights in invention to contractors. The FAR, following the act and the presidential Memorandum, now entails that each contractor may, after required disclosure to the Government, elect to retain title to any subject invention (FAR 27.301)[18]. Subject invention is defined as any invention made, or first reduced to practice, in the performance of work under a Government contract[19]. To retain the title the contractor must

[17]Even if the act was first thought to be addressed only to small business and non-profit firms, the president Memorandum issued by Reagan in 1983 extended its scope to large and for-profit enterprises and, therefore, also the FAR prescriptions on intellectual property right to every entity involved in a contract.

[18]There are some exception to this rule. The Department of Energy and the NASA may retain title to inventions made by the contractor for specific technologies.

[19]Reduction to practice is in turn defined as workable version of the invention created during the performance period. It often occurs after conception. Hence the Government may obtain some rights in already existing conceptions (sometimes even

notify the government the discovery of a patentable invention, and then timely file a patent application. If the contractor retains ownership of the invention the FAR requires that the Government shall have a non-exclusive, irrevocable, paid-up (i.e. no royalties) license to use the invention or to have someone else use the invention on its behalf (FAR 27.302). The rational here is clearly to avoid the government paying twice for the same invention. The most important requirement in the context of this work is that, in order to legally ensure the paid-up license to the government, the FAR obligates the contractor to include into the patent document a government interest statement, reporting that the invention was made with Government support,and that the government has certain rights in the invention (FAR 52.227-11). This rule applies to all procurement contracts that involve the performance of experimental, developmental, or research work, therefore to all procurement contracts thought to produce new knowledge and innovation.

Using the USPTO Full-text and Image Database described in the previous section, it is possible to identify those patents that includes the government interest statement, and also to disentangle between patents that include the statement but originated from a grant by the federal government, and those patents derived from a contract with the government. Among the 32,323 patents that I earlier defined as potentially related to public procurement I hence select those that include the government interest statement and refer to a contract and not to a grant. In this way I identify 1,029 patents, that I deem as 'patent related to public procurement' since they i)belong to a firm that won at least one procurement contract in FY 2000 (97 different assignees); have the priority date in 1999 or 2000; iii) include the government interest statement[20]. Table 3 reports the distribution of these patents across macro-technological field, showing how most of these patents come from electronics and ICT, but also that chemicals and mechanicals account for a considerable share.

[Table 3 about here.]

[Table 4 about here.]

Table 4 report instead some descriptives for patents' and assignees' characteristics for procurement related patents. On average patents in this group make 10.5 backward citations (10,620 cites in total), makes 21.8 claims, and in almost every case are assigned to U.S. contractors. Contractors have on average received 466 million dollars through public procurement contracts. Though the distribution of this measure is highly skewed, with a single contractor obtaining more than 16 billion dollars (Lockheed Martin Corporation), more than half of the firm obtained more than 4,5 million dollars in public procurement contracts. It should also be noted that I do not claim that the 1,029 patents identified with the described strategy represent the whole universe of innovations related to federal public procurement contracts in Fiscal Year 2000. There are in fact few cases in which this strategy would not work: a firm can opt for keeping an invention secret, rather than filing a patent application for it; as we mentioned, some federal agency may retain title on specific technologies and, moreover, when the disclosure of an invention might be detrimental for the national security, the Government may withhold patent application, imposing a secrecy order; a firm could intentionally fail to report innovation obtained in the performance of work under a government contract and hence avoid to include the government interest statement in the patent document[21]. Nonetheless, this strategy ensures that the identified patents are undoubtedly related to innovative public procurement contracts and it therefore allows me to build my treatment variable.

In order to define the treatment variable I check if any of the 10,620 citations done by the 1,029 patent related to public procurement (FY 2000) is going to focal patents in the sample described in section 4.2.1. I

patent pending inventions) due to its involvement in the development of the first working prototype (?).

[20]Sampat and Lichtenberg (2011), Rai and Sampat (2012), and Azoulay et al. (2013) recently implemented a similar strategy to spot patent resulting from grants of the National Institute of Health (NIH).

[21]A report on this topic by the Government Accountability Office conducted in 1999 (GAO, 1999) highlighted some discrepancies between the number of patents that include the government interest statement and the ones the government actually is aware of having rights on. Moreover it also suggested that, in a non-negligible number of cases (from 10 to 20 per cent of the patents in their sample), grantees and contractors failed to add the government interest statement in the patent document, even if it had to be included.

hence define the variable *Treatment_Procurement* to take the value 1 if a patent in the sample (application year 93-97, more than 10 cites received, belonging to a public company) receives a forward citation by one of the 1,029 patents related to public procurement, and the value 0 if such a citation does not arrive. In this way I am able to identify the 903 patents that constitute my treatment group, while the other 70,539 are considered as not treated and will be potential candidates for the control group.

In section 4.4 I will present accurate descriptive statistics for the two groups to carefully analyze similarities and difference between them. Moreover, to avoid potential confounding effects, we also eliminate both from the treated and not treaded groups those patents that receive a citation from subsequents patent that included the government interest statement and whose priority date is before year 1999. In this way I am sure that the treatment related citation is the first citation from a patent related to public procurement for the focal patents.

4.2.3. The outcome variable: the generality index

Since I am interested in evaluating the impact of the treatment on the change in the degree of generality of a patent across time between the treated and the control group of patents, I follow Trajtenberg et al. (1997) to build the outcome variable of my quasi-experiment. As briefly described in section 3.1, Trajtenberg et al. (1997) suggested to look at how forward citations are spread across different technological fields (proxied by patent class), to compute a measure of technological pervasiveness. In particular they develop a Generality index, measured at the single patent level, that is defined as:

$$G_i = (1 - \sum_{j=1}^{J} N_{ij}^2 / N_i) \tag{1}$$

where Nij is the number of forward citations received by patent i from patents in technological class j, while Ni is instead the total number of forward citations received from a patent. The summation term is therefore the Herfindahl concentration index and reports the degree of concentration of forward citations across patent classes. Being one minus the Herfindahl index, the Generality index is also bounded between 0 and 1. In particular, it will get closer to 1 as a patent receives citations from patents belonging to many different patent classes, while it approximate 0 as its forward citation are concentrated in a few classes. Hall (2005) noted that the Generality index proposed by Trajtenberg et al. (1997) suffered of a bias due to the count nature of citation data, a bias that could be particularly important in the case of a low number of forward citations received (N_i). Hall therefore proposed a corrected version of the Generality index, which defines as:

$$\Gamma i = (N_i / N_i - 1)(G_i) \tag{2}$$

where N_i is still the total number of forward citations received by a patent and G_i is the Generality index as defined above. Since my aim is to identify the effect of receiving a citation from a patent related to public procurement on the change over time of a patent's generality level, I measure the adjusted generality index Γ_{it} for all the patents in my sample at two different points in time: the first time as of the beginning of 1999 and the second in the last available year, 2006. We hence compute the number of citations obtained by each patent until the start of year 1999 , their concentration level across USPTO patent class, measured at the 3-digit level (i.e. 400 different technological fields), and then derive the generality index Γ_{i99} for each patent i as of January 1$^{\text{st}}$ 1999. I then compute the Generality index Γ_{i06}, for each patent, at the end of the period for which we have patent and citation data available, December 31$^{\text{st}}$ 2006.

4.3. Empirical approach: CDiD

The main hypothesis of this work is that a patent receiving a citation from a subsequent patent related to a public procurement contract will have a higher generality index with respect to the counter factual situation in which such reference does not arrive. Since the counter factual situation is clearly not observable, I design the quasi-experiment portrayed in figure 5, to recover the average treatment effect exploiting the information coming from not treated patents to estimate it.

[Figure 5 about here.]

I hence select a sample of patents whose application dates lies in the 5-year temporal window from 1993 to 1997. I measure their Generality index as of as of January 1st 1999, Γ_{i99}, on the basis of the citations they obtained up to that moment. In year 1999 or 2000 the treatment arrives: 903 patents receive a forward citation from a patent related to public procurement and I therefore consider them as treated. 70,539 patent do not receive such a citation and are hence considered as not treated. The generality index, Γ_{i06}, is measured for every patent once again at the end of the period for which we have data available, December 31st 2006. Since I hypothesize that generality is a dynamic characteristic that can be, at least to some extent, cultivated over time, in order to corroborate my hypothesis on the role of public procurement in this nurturing process, I will look at the average difference in the outcome variable Γ_i between year 2006 and 1999, and compare the change in generality for treated patents with respect to the control ones. However If I used all the patents in the not treated group to estimate the average treatment effect, results might be biased due to multiple source of potential endogeneity, mainly due to selection bias. In that case I would in fact recover the difference-in differences estimator (DiD), which is formally defined as:

$$DiD = [E(Y_{t1}^{T}|T=1) - E(Y_{t0}^{T}|T=0)] - [E(Y_{t1}^{C}|T=0) - E(Y_{t0}^{C}|T=0)] \tag{3}$$

where $Y_{t1}|T$ and $Y_{t0}|T$ is the outcome variable (Γ_i, in our context) measure at time 1 (2006) and time 0 (1999) for the treated group, while $Y_{t1}|C$ and $Y_{t1}|C$ is the outcome variable at time 1 and time 0 for the not treated group. The main idea is to correct the simple difference before and after the treatment for the treated group, subtracting the simple difference for the not treated group. The estimate of the average treatment effect provided by the DiD estimator is unbiased if and only if the "parallel trend assumption" holds. This assumption involves that, in the absence of the treatment, the average change in Y_{t1} - Y_{t0} for the treated and the control group should have been equal, and hence that, without the treatment, the DiD estimator in equation 3 should be clearly equal to 0. The parallel trend assumption is most likely implausible if pre-treatment characteristics are unbalanced between the treated and the untreated group and they may interact with the dynamics of the outcome variable (Abadie, 2005). In that case the treated and the control group would behave in a different way even in the absence of the treatment and we would hence incur into selection bias, i.e.:

$$[E(Y_{t1}^{T}|T=0) - E(Y_{t0}^{T}|T=0)] - [E(Y_{t1}^{C}|T=0) - E(Y_{t0}^{C}|T=0)] \neq 0 \tag{4}$$

In this context selection bias may arise from the fact that the patents that receive a citation from public procurement related patents might be intrinsically different from the one who do not receive such a citation. In particular, this sort of bias may have a dual origin. On the one side, public procurement related patents, being patents that arises in a particular situation, may present some singularities that lead them to cite patents with the same (or other) specific attributes, making our treated patents different from the non-treated ones. For instance treated patents may belong to a specific subset of patent classes, or they might be more original than non-treated ones. On the other side, specific patents may posses some peculiar feature that rises their probability of receiving the treatment, i.e. a citation from a patent related to public procurement. For example some patents may be more general, more important, or have higher quality, regardless of the treatment, and these features could increase their probability of receiving any kind of citation, included the ones from patents related to public procurement.

In estimating the average treatment effect we have therefore to consider that the whole group of non-treated patents cannot be directly used as proxy for the counterfactual situation in which treated patents do not receive a citation from a patent related to public procurement. To mitigate the potential selection bias I here follow the idea implemented in Lanjouw and Schankerman (2004), Feldman and Yoon (2012), Fier and Pyka (2012), and Czarnitzki et al. (2011), to construct a control group of patents similar to the treated ones along several dimensions. Exploiting the abundant information about patents and patents' assignee in our dataset and the fact that our dependent variable is measured at two points in time, I implement the conditional difference in differences (CDiD) approach, first introduced by Heckman et al. (1998). CDiD, combining the advantages of non-parametric matching method to the ones of the common DiD strategy, allows to tackle the selection on observables and the selection on unobservable issue at the same time. As a first step, it involves in fact the matching of each treated units with a suitable control on the basis of their predicted probability of being treated , i.e. propensity score matching. In a second step, to remove the problem of selection on unobservables, it encompasses the classical before and after comparison (DiD) between the treated and the control group, but only for the matched samples.

4.3.1. Propensity score matching

The main idea of non-paramentric matching method is to find a group of non-treated individuals that are similar to the treated ones in all the relevant pre-treatment characteristics, and to use this group as a close substitute for the unobservable counterfactual situation in which the treated group is not receiving the treatment (Caliendo and Kopeinig, 2008). For a consistent estimation of the treatment effect through matching two conditions have to hold. The first one is the conditional independence assumption (CIA). It requires the assignment to treatment to be independent from the outcome, conditional on a set of observable covariates (X). Rosenbaum and Rubin (1983) show that it is possible to summarize the vector of relevant covariates into a single scalar index, the propensity score, which is defined as the probability of being treated conditional to observable and relevant pre-treatment characteristics. The CIA then formally states :

$$(Y^C; Y^T) \perp D|P(X) \tag{5}$$

where D is the assignment to treatment, and *P(X)* is the probability of receiving the treatment given the relevant covariates.

The second condition that has to hold is the common support condition, formally:

$$0 < P(T|X) < 1 \tag{6}$$

It requires that the relevant observable characteristics are not able to perfectly predict whether a unit is assigned to the treated or to the control group and, therefore, that units sharing the same pre-treatment attributes can be found both in the treated or in the control group with positive probabilities. If both conditions hold, the treated and the control group, once matched on the basis of the propensity score, should be on average observationally identical.

4.3.2. Conditional difference in differences

Though the propensity score matching rules out the selection on observables problem, unobservable characteristics of different units may still affect their probability to receive the treatment, biasing the results. In the context of my analysis permanent unobserved heterogeneity may stem from patents' unobserved attributes such as the inventors' characteristics. In order to allow for time-invariant unobserved heterogeneity, I then implement the before-after comparison of the outcome variable on the matched sample, recovering the average treatment effect thorough the conditional difference in difference estimator (CDiD), as proposed by Heckman et al. (1998):

$$CDiD = E(Y^T_{t1} - Y^T_{t0}, D = 1) - E(Y^C_{t1} - Y^C_{t0}|P(X), D = 0) \tag{7}$$

4.4. Descriptives

Table 5 and 6 report descriptive statistics for the treated and the whole group of non-treated patent both for patents' and assignees' characteristics.

[Table 5 about here.]

[Table 6 about here.]

As hypothesized in section 4.3, selection bias proves to be a very relevant issue in the context of this evaluation. In the first place, table 6 shows that the distribution across technological fields (even when measured via the 6 macro HJT category) presents some dissimilarities between treated and non-treated patents. The latter shows a higher concentration in the Computer and Communications and Drugs and Medicals fields with respect to the former, while the opposite is true for the Chemicals, Electrical and Electronics, and Mechanicals.

Table 6 also reports that patents receiving a forward citation from procurement related patents, i.e. treated patents, are on average substantially more original, receive more citations and tend to have an higher generality index Γ, both before and after the treatment. Moreover, treated patents seems to belong to firms with distinct characteristics. 85 per cent of the treated patents belongs to U.S. corporation while among the non-treated only 75 per cent of the assignees is a U.S. firm. Firms owning treated patents are on average larger both in terms of employees and sales, and to invest more in R&D in absolute terms.

In order to account for these dissimilarities, in evaluating the average treatment effect I have to construct an adequate control group for the treated units and I hence implement the CDiD strategy described in section 4.3. In the next section I will illustrate the first step of this procedure, the propensity score matching, while in the following one I will eventually present the results retrieved by the CDiD estimator.

4.5. Propensity score matching

4.5.1. Propensity score specification and estimation

As mentioned in section 4.3, the propensity score is a measure of the probability for an unit to be treated, conditional to a set of relevant characteristics. The first step to recover the propensity score is hence the detection of the relevant variables affecting the probability of receiving the treatment. Caliendo and Kopeinig (2008) recommend to include only variables that influence simultaneously the treatment participation and the outcome variable, but that are not themselves affected by the treatment, better if time-invariant or measured before the treatment. The decision of which variables to include should hence be taken on the ground of economic theory and previous research (Caliendo and Kopeinig, 2008).

In the first place, I therefore make reference to previous works that used treatment models in the context of a patent analysis and in particular to Lanjouw and Schankerman (2004), Feldman and Yoon (2012), Fier and Pyka (2012), and Czarnitzki et al. (2011). Secondly I consider the peculiar nature of my treatment variable, receiving a citation from a public procurement related patent in year 1999-2000, and evaluate which variables will affect the probability of the arrival of this particular citation and of new citations in general. In all of the afore mentioned works several key characteristics of a patent are taken into account such as its timing, the technological field, the origin of the applicant. Moreover, they include different variables to account for the scope, the originality, and the importance of the patent. These latter attributes are particularly relevant in our case since the economic literature on patents suggests that several patent characteristics give information about its quality and are hence correlated with the number of forward citations that it may receive. For these reasons, among the variables I will use to estimate the propensity score, I primarily include binary variables related to the time of filing of the patent, the *Application Year*, and the *Application-Grant Lag*, which measure the number of years that passed between the filing date and the granting of the patent. It is in fact possible that patents issued earlier (or later) have a different odd of being cited by a procurement related patent in year 1999-2000. In the same way patents granted later may have been around for less time with respect to patents filed in the same year but granted earlier and might be hence less visible and less cited. Given that our dependent variable is based on patent citations, matching on the application year and the application-grant lag also ensures to remove problems related to the truncation bias of patent citations (Fier and Pyka, 2012). As also the descriptive showed, patents in given technological fields may have higher chances of receiving the treatment related citation, hence I include one dummy variable for each main *USPTO Patent class* (3-digit).

Since the scope of the patent is thought to be correlated with the number of citations that a patent receives (Lanjouw and Schankerman, 1999) and also with its economic value (Lerner, 1994), I add the variable *Number of claims* to proxy for the width of the monopoly power granted to a patent (as suggested by Hall et al. (2001)). Some scholars suggest that also the number of backward citations made by a patent can be used to proxy the patent scope (Harhoff et al., 2003), others put forward that instead the number of backward references may measure the crowdedness of a given technological area (Lanjouw and Schankerman, 2001), all in all empirical evidence tends to support the idea that backward citations have a positive correlation with the number of forward citations received by a patent. I hence include the variable *Citations made* among the relevant covariates used to estimate the propensity score. Since not only the number of backward citations, but also their distribution among different technological fields, being a proxy for the basicness of a patent (Trajtenberg et al., 1997), correlates both with the arrival of forward citations and with its technological impact, I also add the variable *Originality*, described in section 4.2.3, to account for the degree of novelty of the patent. I also include the variable *Number of citation 1999* and Γ_{1999}, to consider the technological and economic impact of the patent right before the arrival of the treatment.

I cannot include the variable *Number of Citations 2006* in the vector of relevant covariates, since it is clearly affected by the receipt of the treatment and it is measured at the end of the period we are taking into account. Nevertheless to account for potential systematic difference across treated and non-treated patents along this dimension, I will add the difference in the number of citations received by a specific patent between 1999 and 2006 as a control variable in the in CDiD regression.
In addition to variables collecting information about patent's attributes I also include variables that grasp information about the firm who first applied for the patent, recorded in the year of the filing. The rational for this choice is twofold. On the one side, patents belonging to larger and well performing firms might be more visible and therefore cited more often. On the other, firms investing higher share of their revenues in R&D might also engage more resources in basic research, and they could hence have higher chances to introduce major technological breakthrough. I then include the variables *U.S.Corporation*, *Size*, *Sales*, *Net Income*, and *R&D investment* to account for assignees' heterogeneity.

Once that I identified the relevant variables affecting the probability of receiving the treatment, I proceed to the estimation of the propensity score. In order to recover it I run a probit regression of the treatment variable, *Treatment Procurement*, on the set of relevant covariates listed in the previous section and then predict the propensity propensity score. Table 7 reports the results of the regression.

[Table 7 about here.]

As expected several patent and firm characteristics affects the probability of receiving the treatment. For what concerns patents' attributes in particular the technological field, the originality of the patent, the number of citation received until 1999, and the number of claims appear to affect positively the probability of patent to be treated. Unexpectedly the number of citations made by a patent has a negative effect on the likelihood of receiving the treatment, somehow in accordance with the hypothesis by Lanjouw and Schankerman (2001) that a large number of backward citations may characterize more incremental inventions. While the application year dummies do not seem to have any effect on the probability of receiving a citation related to public procurement, the lag between the filing and the granting of the patent have a positive effect, though the p-value is slightly above .1. For what concern the assignees' attributes, being a U.S. based corporation, having large revenues and net income is increasing the treatment likelihood, while, somewhat surprisingly, R&D investment is reducing it.

4.5.2. Matching quality

Once that I have estimated the propensity score, I use it to implement the non-parametric matching[22]. Among the different algorithm available to perform the matching, I implement the nearest-neighbor algorithm, using the information from up to five neighbor and setting a 'caliper' threshold. As Caliendo and Kopeinig (2008) illustrate, the choice of the algorithm to use is a matter of trade-off between bias and efficiency. Using up to 5 control units to proxy for the counterfactual situation allows me to gain efficiency in the estimation, while the caliper threshold, which imposes a tolerance level on the maximum propensity score distance, ensures to reduce potential bias avoiding bad matches[23].

[22] In order to perform the matching I use the stata module psmatch2, developed by Leuven and Sianesi (2003)

[23] As suggested by the rule of thumb first introduced by Rosenbaum and Rubin (1985) I set the caliper option to .02, a value that corresponds approximately to .25 times the standard deviation of the propensity scores recovered with the probit regression.

Before implementing the CDiD estimate on the paired sample, I have to check whether the propensity score matching procedure allows me to consistently estimate the average treatment effect by taking differences in averages between the treated and the matched counterfactuals.
In the first place I have to check whether the common support condition holds. As mentioned in section 4.3, this condition ensures that we estimate only effects in regions where two observations, the former belonging to the treated and the latter to the control group, can have a similar participation probability. Lechner and Gallen (2001) puts forward that it is possible to assess the overlapping between subsamples through a graphic analysis of the propensity score density distribution, for the treated and the control group, before the matching. Figure 6 displays the kernel density distribution for the treated and the control group, before the matching. As the figure shows, though the shape of distributions differs, there is a large overlap between the distribution of the propensity score of the treated and the not treated group, certifying the common support condition to hold.

[Figure 6 about here.]

Secondly I have to find out if the matching on the propensity score actually manages to to balance the distribution of the relevant variables in the control and the treatment group. The literature suggests different methods to evaluate the matching quality. A common methodology, first introduced by Rosenbaum and Rubin (1985), is the two-sample t-test to check for significant differences in covariate means, for both groups, before and after the matching. Table 8 reports the t-test for all the covariates we included in the probit regression to estimate the propensity score for the unmatched and the matched sample.

[Table 8 about here.]

As expected, before the matching, there is significant difference in the mean between the treated and the control group for several variables such as the originality of the patent, the number of Citations, the generality level before the treatment, technological fields and firms' characteristics. As the right side of the table shows, after the matching implementation, all these differences are no longer statistically significant, suggesting a good performance of the matching procedure in balancing the covariates. Rosenbaum and Rubin (1985) also propose to compute the standardized bias and to compare its size before and after the matching, in order to asses the size of the bias reduction obtained through the propensity score matching method. Table 9 reports the mean and the median standardized bias, before and after the matching. Though there is no clear threshold under which it is possible to tell the success of the matching procedure with certainty, a bias reduction below 3 or 5 per cent is generally considered as sufficient (Caliendo and Kopeinig, 2008).

[Table 9 about here.]

As the table shows, after the matching both the the mean and the median standardized bias fall below the two per cent level, confirming the high quality of the matching on the propensity score.
Finally, since intuitively matching procedure is implemented to "correct" for the difference in terms of probability to receive the treatment between the treated and the control group, we can again look at the visual representation of the propensity score distributions, and make a comparison before and after the matching. As figure 7 displays, the difference in the kernel density distribution of the estimated propensity scores abundantly reduces with respect to the pre-matching situation (figure 6), and the two distributions almost perfectly overlap, once again suggesting that the propensity score matching procedure is succesfully correcting for the selection on observable issue.

[Figure 7 about here.]

5. Results

5.1. Results of the CDiD estimator

Since the assessment of the matching quality ensures that I paired the group of treatment patents with a suitable group of control patents, the problem of the selection on observables should be solved, and the

treated and the control group should be on average observationally identical. Nevertheless, selection on unobservables still represent a major concern and might bias the estimate. As described in section 4.3, I therefore take advantage of the longitudinal nature of the dataset and recover the Conditional Difference-in Differences estimator (CDiD), taking the differences in the mean generality level of the generality index Γ, between the treated and control group, over time (i.e. 1999-2000). In particular I use fixed effect regression to recover the CDiD estimator. This allows to control for patent and time fixed effect, eliminating the selection due to time-invariant individual heterogeneity. As briefly discussed in section 4.5, I also control for the difference in the number of citations obtained by a patent in year 1999 and in 2006, that I could not take into account in the propensity score estimation, including the variable *Number of citations*.
Since both the selection on observables and unobservables problem are taken into account, the CDiD estimator consistently identifies the average treatment effect on the treated. In our case it hence identifies the average effect of receiving a citation from a patent related to public procurement (*Treatment_Procurement*) on the change in the generality level Γ for patents that actually receive such a citation.

[Table 10 about here.]

Table 10 presents the results retrieved through the CDiD estimator. As the table shows receiving a citation from a patent related to innovative public procurement has a positive and significant impact on the generality level as measured by the index Γ, confirming the main hypothesis stated in section 3. In particular, receiving the a citation from an innovation complementarity stimulated by public procurement and covered by a patent, on average raises the generality level of the focal patent (the upstream technology) of 3.6 %, compared to the counterfactual situation in which that specific citation did not arrive. As the table 10 shows, the change in the number of citations over time appears to have very small, but still significant, negative effect on the average change in the generality level of the focal patents. This suggests that a new citation is more likely to arrive from a patent class that is already citing the focal patent.

5.2. Robustness checks

5.2.1. The net generality index

The latter consideration may lead to think that the positive and significant impact of the treatment presented in the previous section might be driven only by a direct effect of the citation coming from the procurement related patent. The procurement related citation is in fact included among the citations used to compute the generality index Γ. It might hence be that this specific citation, being the only one arriving from a particular patent class, is directly lowering the concentration index without inducing any change in the knowledge diffusion process with respect to the control patents. In order to rule out this hypothesis, I also compute a second generality index, that I label *Net_Γ*, which is computed exactly in the same way as the previous one (Γ,see section 4.2.3), but removing the specific citation coming from the public procurement related patent. I then compute the CDiD estimator on the matched sample I used in the previous section, adopting *Net_Γ* as the outcome variable. Table 11 reports the results of this robustness check.

[Table 11 about here.]

Results show that the effect of receiving a citation from a patent related to public procurement is still positive and significant even when the procurement related reference is not taken into account for computing the generality index. Even though the magnitude of this results can not be trusted since it is obtained removing systematically one cites from the treated group only, it is worth to note that the size of the impact of the treatment on the average generality change is even larger in this case (close to 5.8 per cent compared with 3.6 per cent in the previous estimation). This means that the afore mentioned hypothesis about a direct effect of the procurement related citations on the average change in the generality index have to be rejected. Removing that specific citation is in fact raising the effect of the treatment and hence decreasing the level of concentration of citations across different patent classes for treated patents. This means, in turns, that on average focal patents receive other citations from patents belonging to the same technological class of the procurement related patents. The direction and the significance of the average treatment effect presented in the previous section appear hence to be robust.

5.2.2. Different measure of generality

A well known drawback of the generality measure developed by Trajtenberg et al. (1997) is that it assumes that all the categories taken into account, i.e. US patent classes, are equidistant from each other in the technology space, while this is clearly not the case in the real world. In order to partially correct for this problem Hall and Trajtenberg (2004) computed different version of the generality index described in equation **??**, using different classification system[24] Following their example I also calculate two additional generality indexes, *Generality_IPC* and *Generality_HJT*, based on two alternative classifications: the main international patent class (IPC, approximately 1200 cells), and the Hall-Jaffe-Trajtenberg technology subcategories (HJT, 36 cells). As stressed by Hall and Trajtenberg (2004), each of these categorization may help to mitigate the technological distance-issue given that on the one side the IPC classification, having more classes (1200), is much more detailed than the USPTO one, and, on the other, the HJT categorization (first developed by (Hall et al., 2001) is based on more equal groups of technologies. Table 12 presents the results for the CDiD estimator computed using the *Generality_IPC* and *Generality_HJT* as outcome variables. As the table shows in both cases the average effect of receiving a citation from a patent related to public procurement is positive, significant and even larger in terms of magnitude with respect to the estimation retrieved using the USPTO classification. This confirms that the result presented in section 5.1 is not driven by the specific classification used to compute the generality index Γ, but is robust to more detailed and balanced categorizations.

[Table 12 about here.]

5.2.3. A placebo test

One common strategy used to validate estimates obtained through the difference in differences approach is to implement a placebo test. This test consists in building a 'fake' treatment group and re-estimate the coefficient for the DiD estimator using the same control group. Obviously since the false treatment group is not receiving the treatment, a positive (negative) significant coefficient for the DiD estimator in the placebo test would suggest that the result of the focal DiD is biased. The control group and the treated group would in fact follow different trends even in the absence of the treatment.

In order to rule out this possibility, I run a placebo test identifying a false treatment group of patents that did not receive a citation from patents related to public procurement but that are similar to the one in the original treated group along the same characteristics I used to build the control group. In particular I use the same strategy that I adopted to find a suitable control group to spot credible 'fake-treated' patents. Clearly, patents in the control group are not available to be selected as patent in the fake treatment group. I hence rerun the propensity score matching method as done (through probit regression and nearest neighbor algorithm) in section 4.5. In this way I identify 886 fake-treated patents and I then recover the DiD estimator comparing them with same control group used in the original estimation.

[Table 13 about here.]

Table 13 reports the results of the placebo CDiD. As it is possible to see, the effect of the fake-treatment is close to zero and is not statistically significant. The placebo CDiD result therefore confirms that the finding of the original CDiD are not driven by some underlying difference in trends between the treated and the control, but only by the actual treatment.

5.2.4. Matching only on patent characteristics

Finally, I run an additional robustness check in which I use only patent-specific characteristics (i.e. not assignees' attributes) to find suitable control patents for the treated units. The rational for performing this test is twofold. In the first place, I would like to rule out the concern that the results showed in the previous

[24]In particular they use: US patent class (approximately 400 cells), Hall-Jaffe-Trajtenberg technology subcategories (36 cells), Main International Patent Class (approximately 1200 cells), Industry classification based on Silvermanâs IPC-SIC concordance (Silverman 2002) for industry of manufacture, aggregated to the Hall-Vopel (1997) level (37 cells), and Industry classification based on Silverman's IPC-SIC concordance for industry of use, aggregated to Hall-Vopel level (37 cells).

sections are driven by the fact that I considered only patents owned by public companies. Secondly, by including patents owned by any kind of assignee, I will be able to enlarge the treated group and therefore also to reduce potential drawbacks due to the limited size of the treated group

In order to perform this robustness check I hence apply the same empirical method described in section 4.3, but without limiting its scope to patents belonging to public companies. The sample then includes patents owned by private and public companies, research institutes, universities, or individuals (coming from U.S. and other countries) and reaches 170,226 units. The whole sample is then split in treated and not treated group exploiting the same strategy used in the previous case. In this way I end up with a treated group composed by 1,524 units, that I match with up to 5 suitable control patents in order to recover the average treatment effect on the treated, again using the same CDiD technique portrayed in sections 4.3 and 4.5[25].

[Table 14 about here.]

only Table 14 reports the result of CDiD estimators for the group matched only on patent characteristics. As the table shows receiving a citation from a patent related to public procurement still has a positive and significant effect on the degree of generality of the cited patent. Moreover, the magnitude of the impact seems to be larger with respect to the case in which only public companies were considered. Even though the size of the effect cannot be entirely trusted since, as showed in section4.5, assignee specific attributes play an important role in determining treatment participation, this result confirm that our the estimates retrieved in the focal CDiD were not driven by the fact that only public companies' patents were taken into account.

6. Conclusions

Even though the contributions of economic historians and the theoretical work in the GPT literature have long suggested a tight relation between public demand and the emergence of radical innovation, no econometric work tried to investigate the role of innovative public procurement in raising the technological generality of given technologies. In this paper I aimed at filling this gap through a patent data analysis.

Even if the literature in some case suggested that government demand has been able to directly produce major technological breakthrough, here I hypothesized that public procurement, by stimulating additional innovation complementarities in the application sectors, can increase the likelihood of diffusion of upstream technologies among different sectors, making them more pervasive with respect to the counterfactual situation in which no stimulus from public demand occurred. To empirically test this hypothesis I exploited patent data and in particular patent citations. Citations in fact allowed me to identify the linkages between innovations induced by public procurement and their technological antecedents, and also to measure the degree of pervasiveness of patents through the Generality Index, first introduced by Trajtenberg et al. (1997). On this ground I formally hypothesized that receiving a citation from a patent related to public procurement raises the generality level of the cited patent.

I hence designed a quasi-experiment exploiting the information of the original dataset I created using data from four different sources: i) NBER patent data project; ii) Federal Procurement Data System (FPDS); iii) USPTO full text and images database; iv) the Compustat North America Database. I compared the change in the generality level Γ at two different points in time, 1999 and 2006, between treated and a control patents, whose application date falls in the period 1993-1997 and who received at least 10 forward citations in 2006. A patent was put into the treatment group if it received a citation from a patent related to public procurement in year 1999 or 2000. In order to mitigate the potential bias due to selection on observed and unobserved variables, I implemented a conditional difference-in-differences approach, matching treated patents with a group of suitable control patents and looking at the change in generality over time. The results of the paper suggest a positive and significant impact of innovative public procurement upon the generality of upstream

[25]The only difference is clearly that I do not include variables that grasp assignees' specific characteristics in the probit regression used to estimate the propensity score, but only patent attributes

technologies. In particular, on average receiving a citation by a patent related to public procurement raises the Generality Index of the cited patent of 3.6%, confirming the initial hypothesis.

As stressed throughout the paper I did not put forward that the mere presence of public procurement in the application sectors can always lead to the arrival of very pervasive or radical technologies but only that it may create the right soil to 'cultivate' technologies that may, or may not, have the potential to become very general. On this ground, the results of this paper appear to provide empirical support to the idea that public demand could play a crucial role in setting in motion the virtuous cycle that may lead to, or simply accelerate, the deployment of a new general purpose technology. 'Schumpeterian demand policies' that pay attention to the technological composition of public procurement might then represent an effective policy tool to spur innovation bandwagons and radical technological change (Antonelli, 2010). Moreover, this technology intensive public demand might be mostly useful in those fields,, like green technologies, where the private sector is not willing to invest enough due to the high degree of uncertainty (Mazzucato, 2013).

References

Abadie, A. (2005). Semiparametric difference-in-differences estimators. *The Review of Economic Studies*, 72(1):1–19.

Antonelli, C. (2010). *La mossa del cavallo*. Rosenberg & Sellier.

Aschhoff, B. and Sofka, W. (2009). Innovation on demand. can public procurement drive market success of innovations? *Research Policy*, 38(8):1235 – 1247.

Azoulay, P., Graff-Zivin, J., Li, D., and Sampat, B. (2013). Public r&d investments and private sector patenting: Evidence from nih funding rules.

Bewley, T. F. (2001). Market innovation and entrepreneurship: A knightian view. *Economics Essays: A Festschrift for Werner Hildenbrand*, page 41.

Bloch, D. and Gray, R. (2012). *Intellectual property in government contracts: protecting and enforcing IP at the state and federal level*. Oxford University Press.

Brammer, S. and Walker, H. (2011). Sustainable procurement in the public sector: an international comparative study. *International Journal of Operations & Production Management*, 31(4):452–476.

Bresnahan, T. and Trajtenberg, M. (1995). General purpose technologies 'Engines of growth'? *Journal of econometrics*, 65:83–108.

Caliendo, M. and Kopeinig, S. (2008). Some practical guidance for the implementation of propensity score matching. *Journal of Economic Surveys*, 22(1):31–72.

Cantner, U. and Vannuccini, S. (2012). A new view of general purpose technologies.

Carlaw, K. I. and Lipsey, R. G. (2011). Sustained endogenous growth driven by structured and evolving general purpose technologies. *Journal of Evolutionary Economics*, 21(4):563–593.

Cowan, R. and Foray, D. (1995). Quandaries in the economics of dual technologies and spillovers from military to civilian research and development. *Research Policy*, 24(6):851–868.

Cozzi, G. and Impullitti, G. (2010). Government spending composition, technical change, and wage inequality. *Journal of the European Economic Association*, 8(6):1325–1358.

Czarnitzki, D., Hussinger, K., and Schneider, C. (2011). Commercializing academic research: the quality of faculty patenting. *Industrial and Corporate Change*, 20(5):1403–1437.

Dalpé, R. (1994). Effects of government procurement on industrial innovation. *Technology in Society*, 16(1):65–83.

David, P. (1989). The dynamo and the computer: an historical perspective on the modern productivity paradox. *The American Economic Review*.

Dosi, G. (1982). Technological paradigms and technological trajectories. *Research Policy*, 11(3):147–162.

Edler, J. and Georghiou, L. (2007). Public procurement and innovation. resurrecting the demand side. *Research Policy*, 36(7):949–963.

Edquist, C. and Hommen, L. (2000a). *Public technology procurement and innovation*, volume 16. Springer.

Edquist, C. and Hommen, L. (2000b). Public technology procurement and innovation theory. *Public technology procurement and innovation*, pages 5–70.

EU (2010). *Europe 2020 Flagship Initiative: Innovation Union*. European Union, Brussels.

Fabrizio, K. R. and Mowery, D. C. (2007). The federal role in financing major innovations: information technology during the postwar period. *Financing innovation in the United States*, pages 283–316.

Feldman, M. and Yoon, J. (2012). An empirical test for general purpose technology: an examination of the cohen–boyer rdna technology. *Industrial and Corporate Change*, 21(2):249–275.

Fier, H. and Pyka, A. (2012). Is it worth all the trouble? an assessment of the economic value of firm patent applications with shared intellectual property rights in the biotechnology industry. In *Technology Transfer in a Global Economy*, pages 123–142. Springer.

Flamm, K. (1987). *Targeting the computer: Government support and international competition*. The Brookings Institution.

Flanagan, K., Uyarra, E., and Laranja, M. (2011). Reconceptualising thepolicy mix for innovation. *Research Policy*.

Fontana, R. and Guerzoni, M. (2008). Incentives and uncertainty: an empirical analysis of the impact of demand on innovation. *Cambridge Journal of Economics*, 32(6):927–946.

Foray, D. (1997). The dynamic implications of increasing returns: Technological change and path dependent inefficiency. *International Journal of Industrial Organization*, 15(6):733–752.

Foray, D., Mowery, D., and Nelson, R. (2012). Public r&d and social challenges: What lessons from mission r&d programs? *Research Policy*, 41(10):1697–1702.

GAO (1999). Reporting requirements for federally sponsored inventions need revision. Technical report, GAO/RCED-99-242.

Geroski, P. (1990). Procurement policy as a tool of industrial policy. *International Review of Applied Economics*, 4(2):182–198.

Geroski, P. A. (2000). Models of technology diffusion. *Research policy*, 29(4):603–625.

Griliches, Z. (1991). Patent statistics as economic indicators: A survey. *Journal of Economic Literature*, 28(4):1661–1707.

Guerzoni, M. and Raiteri, E. (2013). Demand vs. supply technology policies: hidden treatment and new empirical evidence on the technology policy mix. *Bureau of Research in Innovation, Torino*.

Hall, B., Jaffe, A., and Trajtenberg, M. (2006). Market value and patent citations'. *INTERNATIONAL LIBRARY OF ...*, 36(1):16–38.

Hall, B., Jaffe, A. B., and Trajtenberg, M. (2001). The nber patent citation data file: Lessons, insights and methodological tools. (8498).

Hall, B. and Trajtenberg, M. (2004). Uncovering gpts with patent data.

Hall, B. H. (2005). A note on the bias in herfindahl-type measures based on count data. *Revue d'Économie Industrielle*, 110(1):149–156.

Hall, B. H. (2006). Innovation and diffusion. *The Oxford Handbook of Innovation*, page 459.

Harhoff, D., Scherer, F. M., and Vopel, K. (2003). Citations, family size, opposition and the value of patent rights. *Research Policy*, 32(8):1343–1363.

Heckman, J., Ichimura, H., Smith, J., and Todd, P. (1998). Characterizing selection bias using experimental data. *Econometrica*, 66(5):1017–1098.

Helpman, E. and Trajtenberg, M. (1994). A time to sow and a time to reap: Growth based on general purpose technologies.

Henderson, R., Jaffe, A. B., and Trajtenberg, M. (1998). Universities as a source of commercial technology: a detailed analysis of university patenting, 1965-1988. *Review of Economics and Statistics*, 80(1):119–127.

Jaffe, A., Trajtenberg, M., and Fogarty, M. (2000). The meaning of patent citations: report on the nber/case-western reserve survey of patentees.

Jaffe, A., Trajtenberg, M., and Henderson, R. (1993). Geographic localization of knowledge spillovers as evidenced by patent citations. *the Quarterly journal of Economics*, 108(3):577–598.

Kaldor, N. (1966). *Causes of the slow rate of economic growth of the United Kingdom: an inaugural lecture.* Cambridge University Press.

Katz, B. and Phillips, A. (1982). The computer industry. *Government and Technical Progress: A Cross-Industry Analysis*, pages 162–232.

Langlois, R. N. and Steinmueller, W. E. (1999). The evolution of competitive advantage in the worldwide semiconductor industry, 1947-1996. *The Sources of Industrial Leadership*, pages 19–78.

Lanjouw, J. and Schankerman, M. (2004). Patent quality and research productivity: Measuring innovation with multiple indicators*. *The Economic Journal*, 114:441–465.

Lanjouw, J. O. and Schankerman, M. (1999). The quality of ideas: measuring innovation with multiple indicators. Technical report, National bureau of economic research.

Lanjouw, J. O. and Schankerman, M. (2001). Characteristics of patent litigation: a window on competition. *RAND journal of economics*, pages 129–151.

Lechner, M. and Gallen, S. (2001). *A note on the common support problem in applied evaluation studies.* Forschungsgemeinschaft für Nationalökonomie.

Lerner, J. (1994). The importance of patent scope: an empirical analysis. *The RAND Journal of Economics*, pages 319–333.

Leuven, E. and Sianesi, B. (2003). Psmatch2: Stata module to perform full mahalanobis and propensity score matching, common support graphing, and covariate imbalance testing. Statistical Software Components, Boston College Department of Economics.

Levin, R. C. (1982). The semiconductor industry. *Government and Technical Progress: A cross-industry analysis*, 9:76–77.

Lichtenberg, F. R. (1988). The private r and d investment response to federal design and technical competitions. *The American Economic Review*, 78(3):pp. 550–559.

Liebman, J. B. and Mahoney, N. (2013). Do expiring budgets lead to wasteful year-end spending? evidence from federal procurement. Technical report, National Bureau of Economic Research.

Malerba, F., Nelson, R. R., Orsenigo, L., and Winter, S. G. (2007). Demand, innovation, and the dynamics of market structure: The role of experimental users and diverse preferences. *Journal of Evolutionary Economics*, 17(4):371–399.

Mazzucato, M. (2013). *The Entrepreneurial State: Debunking public vs. private sector myths*. Anthem Press.

Mokyr, J. (1990). Punctuated equilibria and technological progress. *The American Economic Review*, pages 350–354.

Moser, P. and Nicholas, T. (2004). Was electricity a general purpose technology? evidence from historical patent citations. *American Economic Review*, pages 388–394.

Mowery, D. and Rosenberg, N. (1989). *Technology and the pursuit of economic growth*. Cambridge University Press, Cambridge.

Mowery, D. C. (2011). Federal policy and the development of semiconductors, computer hardware, and computer software: A policy model for climate change r&d? In *Accelerating Energy Innovation: Insights from Multiple Sectors*, pages 159–188. University of Chicago Press.

Mowery, D. C. (2012). Defense-related r&d as a model for "grand challenges" technology policies. *Research Policy*, 41(10):1703–1715.

Mowery, D. C., Nelson, R. R., and Martin, B. R. (2010). Technology policy and global warming: Why new policy models are needed (or why putting new wine in old bottles won't work). *Research Policy*, 39(8):1011–1023.

Mowery, D. C. and Rosenberg, N. (1979). The influence of market demand upon innovation. a critical review of some recent empirical studies. *Research Policy*, 8(2):102–153.

Mowery, D. C. and Rosenberg, N. (1982). The commercial aircraft industry. *Government and Technological Progress, Pergamon Press, New York*.

Nelson, R. R. (1982). *Government and Technical Progress: a cross-industry analysis*. Pergamon.

Nelson, R. R., Peterhansl, A., and Sampat, B. (2004). Why and how innovations get adopted: a tale of four models. *Industrial and Corporate Change*, 13(5):679–699.

OECD (2013). Government at a glance 2013.

Rai, A. and Sampat, B. (2012). Accountability in patenting of federally funded research. *Nature biotechnology*, 30(10):953.

Rogers, E. (1995). *Diffusion of innovations*. The Free Press.

Rolfstam, M. (2009). Public procurement as an innovation policy tool: the role of institutions. *Science and Public Policy*, 36(5):349–360.

Rolfstam, M. (2012). Understanding public procurement of innovation: Definitions, innovation types and interaction modes. *Available at SSRN: http://ssrn.com/abstract=2011488*.

Rosenbaum, P. and Rubin, D. (1983). The central role of the propensity score in observational studies for causal effects. *Biometrika*, 70(1):41–55.

Rosenbaum, P. R. and Rubin, D. B. (1985). Comments on "on State Education Statistics": A difficulty with regression analyses of regional test score averages. *Journal of Educational Statistics*, 10:326–333.

Rosenberg, N. (1998). Uncertainty and technological change. *Dale Neef, Anthony Siesfeld y Jacquelyn Cefola, The Economic Impact of Knowledge, Boston, Butterworth Heinemann*, pages 17–34.

Rosenberg, N. and Trajtenberg, M. (2004). A general-purpose technology at work: The corliss steam engine in the late-nineteenth-century united states. *The Journal of Economic ...*, (2004).

Ruttan, V. W. (2006). Is war necessary for economic growth? military procurement and technology development. Technical report, University of Minnesota, Department of Applied Economics.

Sampat, B. N. and Lichtenberg, F. R. (2011). What are the respective roles of the public and private sectors in pharmaceutical innovation? *Health Affairs*, 30(2):332–339.

Scherer, F. (1965). Firm size, market structure, opportunity, and the output of patented inventions. *The American Economic Review*, 55(5):1097–1125.

Schmookler, J. (1962). Changes in industry and in the state of knowledge as determinants of industrial invention. *The rate and direction of inventive activity: Economic ...*, pages 195–232.

Schumpeter, J. A. (1934). *The theory of economic development: An inquiry into profits, capital, credit, interest, and the business cycle*, volume 55. Transaction Publishers.

Sharp, G. S. (2003). A layman's guide to intellectual property in defense contracts. *Public Contract Law Journal*, pages 99–137.

Trajtenberg, M. (1990). A penny for your quotes: patent citations and the value of innovations. *The Rand Journal of Economics*, 21(1):172–187.

Trajtenberg, M., Henderson, R., and Jaffe, A. (1997). University versus corporate patents: A window on the basicness of invention. *Economics of Innovation ...*, (September 2012):37–41.

Uyarra, E. (2013). Review of measures in support of public procurement of innovation. *Compendium of Evidence on the Effectiveness of Innovation Policy Intervention*.

Uyarra, E. and Flanagan, K. (2010). Understanding the innovation impacts of public procurement. *European Planning Studies*, 18(1):123–143.

Von Hippel, E. (1988). *The source of innovation*. Oxford University Press, New York.

Youtie, J., Iacopetta, M., and Graham, S. (2008). Assessing the nature of nanotechnology: can we uncover an emerging general purpose technology? *The Journal of Technology Transfer*, 33(3):315–329.

List of Figures

Figure 1: BT diagram

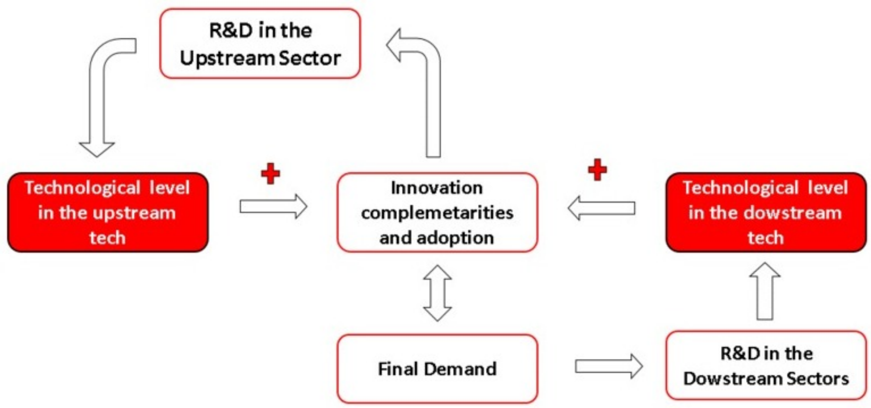

Figure 2: BT diagram with Procurement

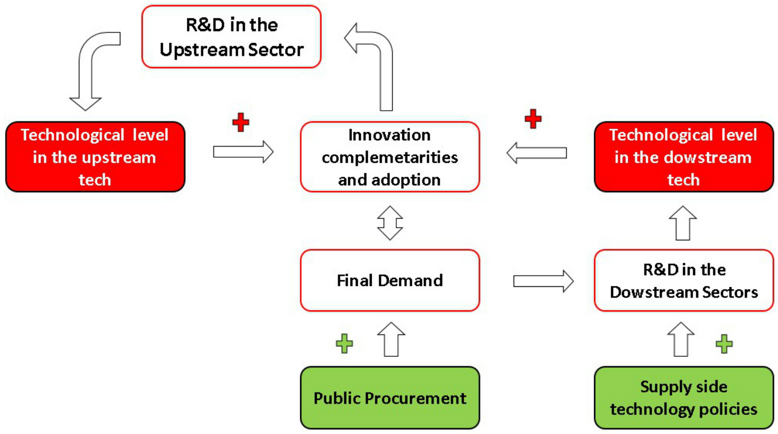

Figure 3: Quasi-experiment

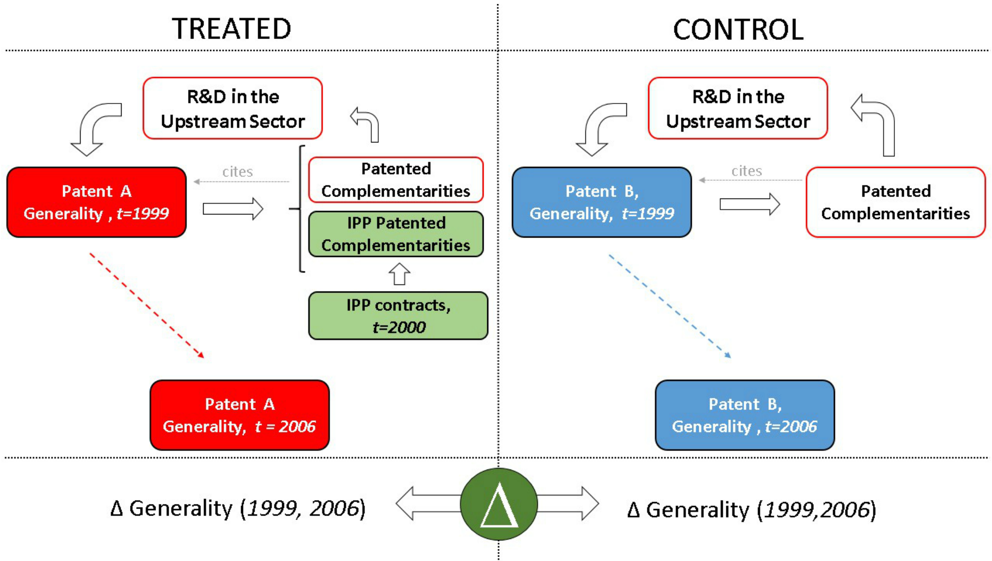

Figure 4: Distribution of Patents according to the number of forward citations received in 2006

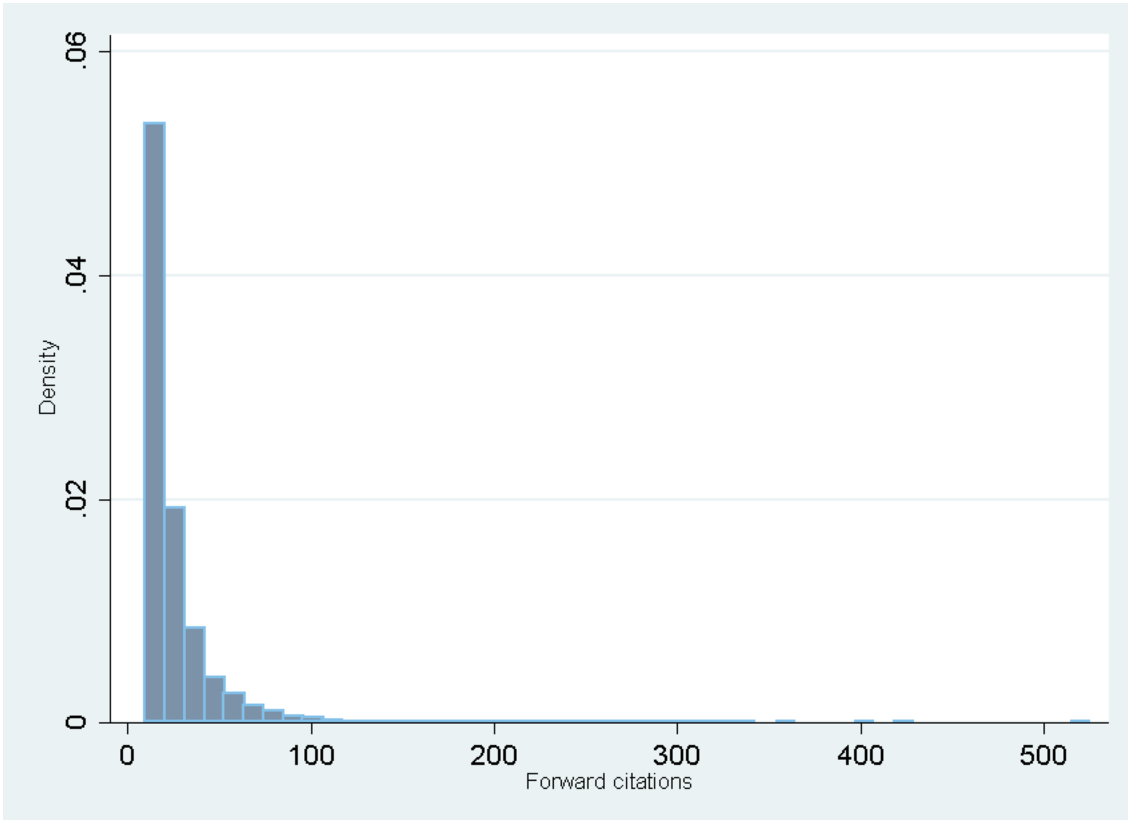

Figure 5: Quasi-experiment

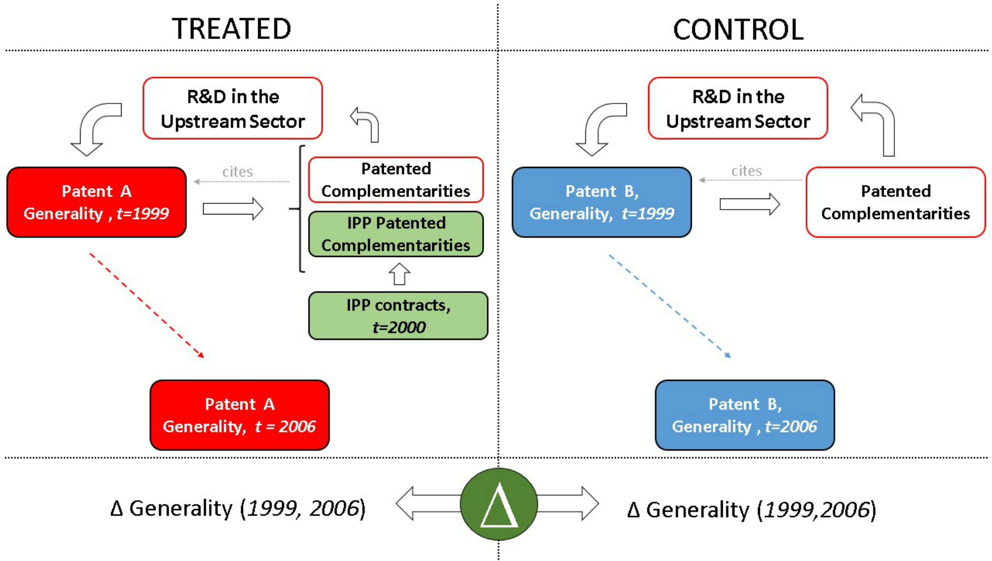

Figure 6: Distributions of the propensity score for treated and not treated group before the matching

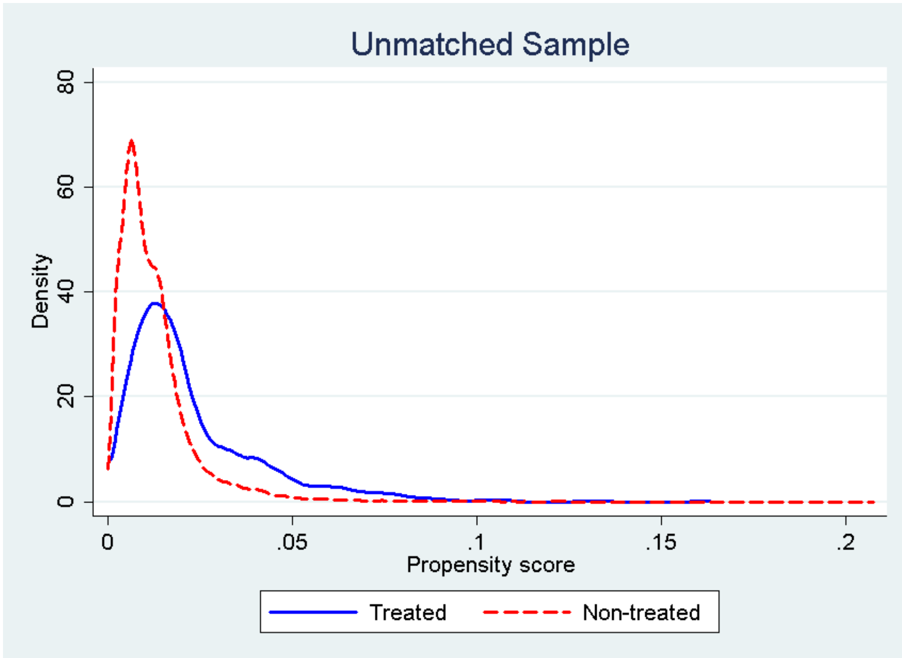

Figure 7: Distributions of the propensity score for treated and not treated group before the matching

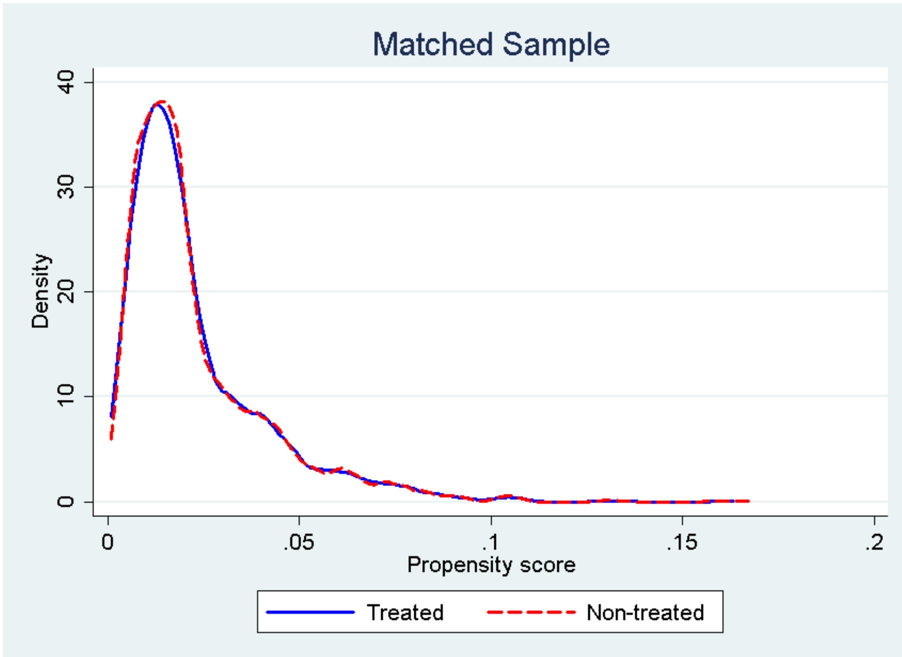

List of Tables

Table 1: Patents distribution across technology class

Category	Technology class(CCL HJT)	N	Share
Chemicals		**7280**	**10.2**
	Agriculture, Food, Textiles	89	0.1
	Coating	901	1.3
	Gas	248	0.3
	Organic Compounds	554	0.8
	Resins	1418	2.0
	Miscellaneous	4070	5.7
Computer and Communications		**29074**	**40.7**
	Communications	9353	13.1
	Hardware and software	8410	11.8
	Peripherals	3788	5.3
	Information storage	6061	8.5
	Miscellaneous	1462	2.0
Drugs and Medical		**6083**	**8.5**
	Drugs	2253	3.2
	Surgery and medical instruments	3455	4.8
	Biotechnology	47	0.1
	Miscellaneous	328	0.5
Electric and Electronics		**17636**	**24.7**
	Electrical devices	2252	3.2
	Electrica Lightning	977	1.4
	Measuring and Testing	1529	2.1
	Nuclear & X-rays	834	1.2
	Power System	3253	4.6
	Semiconductors Devices	6448	9.0
	Miscellaneous	2343	3.3
Mechanical		**6722**	**9.4**
	Material processing & Handling	1130	1.6
	Metal working	927	1.3
	Motors, Engine, Parts	1223	1.7
	Optics	1239	1.7
	Transportation	995	1.4
	Miscellaneous	1208	1.7
Others		**4643**	**6.5**
	Agricolture, Husbandry, Food	197	0.3
	Amusement Devices	158	0.2
	Apparel & Textiles	149	0.2
	Earth Working & Wells	646	0.9
	Furniture	369	0.5
	Heating	166	0.2
	Pipes &Joints	135	0.2
	Receptacles	319	0.4
	Miscellaneous	2504	3.5
Total		**71438**	**100**

Table 2: Descriptives

Patent characteristics	mean	sd
Application-grant lag	1.94	0.83
Cites received 2006	25.36	21.76
Cites made	11.27	13.61
Number of Claims	18.02	13.73
Originality	0.43	0.27
Application Year 1993	0.20	0.40
Application Year 1994	0.22	0.41
Application Year 1995	0.24	0.43
Application Year 1996	0.21	0.41
Application Year 1997	0.14	0.34
Assignee Characteristics	mean	sd
U.S. Corporation	0.75	0.43
Sales	25976	29884
R&D investment	1556	1661
Net income	682	2211
Number of employees	107.8	120.6
N	71438	

Table 3: Distribution across technology class of the patents related to public procurement

Category	**N**	**Share**
Chemicals	195	19.0
Computer and Communications	236	22.9
Drugs and Medical	34	3.3
Electric and Electronics	310	30.1
Mechanicals	155	15.1
Others	99	9.6
Total	**1029**	**100**

Table 4: Descriptives for patents related to public procurement

	mean	sd
Patent characteristics		
number of claims	21.82	16.17
Cites made	10.51	9.54
Application-grant lag	2.49	1.08
Applcation year 1999	0.40	0.49
Applcation year 2000	0.50	0.50
Applcation year 2001	0.08	0.26
Assignee Characteristics		
U.S. entities	.99	0.05
Total procurement dollars (million)	466	1910
Share of R&D procurement	0.46	0.46
Share of DOD procurement	0.56	0.43
Share of competed procurement	0.73	0.33
N	1029	

Table 5: Share of patents across HJT tech category (1-6) by treatment status

HJT tech category	**Share**	
	Non-treated	**Treated**
Chemicals	10.2	12.2
Computer and Communications	40.8	35.8
Drugs and Medical	8.6	3.1
Electric and Electronics	24.6	30.0
Mechanicals	9.4	12.6
Others	6.5	6.3
N	70539	903

Table 6: Descriptives by treatment status

	Non-treated		**Treated**	
	mean	sd	mean	sd
Patent characteristics				
Application-grant lag	1.94	0.83	1.92	0.79
Cites received 2006	25.26	21.51	33.90	35.57
Cites received 1999	9.11	10.14	9.99	14.85
Cites made	11.27	13.65	10.92	10.69
Number of Claims	18.02	13.75	18.69	12.93
Originality	0.43	0.27	0.47	0.27
Γ_2006	0.57	0.23	0.64	0.21
Γ_1999	0.50	0.34	0.53	0.34
Application year 1993	0.20	0.40	0.20	0.40
Application year 1994	0.22	0.41	0.22	0.41
Application year 1995	0.24	0.43	0.23	0.42
Application year 1996	0.21	0.41	0.20	0.40
Application year 1997	0.13	0.34	0.15	0.36
Assignee Characteristics				
U.S. Corporation	0.75	0.44	0.85	0.36
Sales	25858	29806	35125	34240
Net income	673	2199	1404	2915
R&D investment	1555	1660	1705	1764
Number of employees	107.36	120.42	143.25	135.76
N	70539		903	

Table 7: Probit results

	Treated			
	b	se	t	p
Treated				
Number of claims	.0018099*	.0009558	1.893526	.058288
Citations made	-.0023675*	.0012581	-1.881906	.0598487
Originality	.2254436***	.0560177	4.024505	.0000571
Application Year 1994	.0199969	.0414727	.4821694	.6296856
Application Year 1995	-.0220617	.0430022	-.5130373	.6079252
Application Year 1996	-.0115794	.0450699	-.2569209	.7972398
Application Year 1997	.0769461	.0506414	1.51943	.1286543
Application-Grant Lag	.0283305	.0174962	1.619236	.1053965
Number of citation 1999	.0045111***	.0011864	3.802394	.0001433
Γ_{1999}	.0469813	.0425341	1.104557	.2693515
Net Income	.000019***	7.04e-06	2.697946	.0069769
R&D investment	-.0001167***	.0000172	-6.800907	1.04e-11
Sales	7.88e-06***	9.43e-07	8.356067	0
U.S. Corporation	.2325874***	.0372055	6.251425	4.07e-10
HJT subcat2	.2069472*	.1143223	1.810208	.0702635
HJT subcat3	.2234078	.1873057	1.192744	.2329695
HJT subcat4	-.7251165**	.3206278	-2.261553	.0237251
HJT subcat5	-.7389765***	.1822697	-4.054302	.0000503
HJT subcat6	.0943127	.0814678	1.157668	.2469995
HJT subcat7	.0526638	.0763208	.6900324	.4901738
HJT subcat8	-.037311	.0775451	-.4811517	.6304087
HJT subcat9	-.3126514***	.1038059	-3.011883	.0025963
HJT subcat10	-.1396962	.0864068	-1.616727	.1059372
HJT subcat11	-.0853306	.1147693	-.7434969	.4571809
HJT subcat12	-.5456036***	.1499826	-3.637779	.000275
HJT subcat13	-.2502491**	.1021471	-2.44989	.01429
HJT subcat16	-.1988865*	.1135878	-1.750949	.0799546
HJT subcat17	.1310653	.1208216	1.084784	.2780175
HJT subcat18	.3865168***	.0906971	4.261622	.0000203
HJT subcat19	.3919717***	.1076293	3.641866	.0002707
HJT subcat20	.1410234*	.084777	1.663463	.0962197
HJT subcat21	.0424543	.0802936	.5287377	.5969874
HJT subcat22	-.2549342**	.1193446	-2.136119	.0326697
HJT subcat23	-.1624451	.1352674	-1.200918	.2297832
HJT subcat24	.1430589	.1163794	1.229246	.2189795
HJT subcat25	.2193313**	.1024472	2.14092	.0322805
HJT subcat26	.4544197***	.0979239	4.640538	3.48e-06
HJT subcat27	-.5022864***	.1847292	-2.719042	.0065471
HJT subcat28	-.2654577*	.1443956	-1.838406	.0660026
HJT subcat32	-.2152654	.169561	-1.269546	.2042466
HJT subcat34	.601023***	.1732169	3.469771	.0005209
HJT subcat35	.3871127*	.212692	1.820062	.0687495
_cons	-2.705938***	.0931309	-29.0552	0

Table 8: Descriptives statistics for the Unmatched and the Matched Sample

	Not Matched			Matched		
	Treated	Control	T-test m.d.	Treated	Control	T-test m.d.
Number of Claims	18.714	18.018	1.5	18.714	18.687	0.04
Citations Made	10.927	11.183	-0.56	10.927	11.053	-0.24
Originality	.47423	.42818	5.07***	.47423	.48413	-0.80
Application Year 1993	.20222	.19746	0.36	.20222	.20467	-0.13
Application Year 1994	.21778	.21817	-0.03	.21778	.22133	-0.18
Application Year 1995	.22556	.23942	-0.97	.22556	.21881	0.34
Application Year 1996	.20444	.21008	-0.41	.20444	.20889	-0.23
Application Year 1997	.15	.13487	1.32	.15	.1463	0.22
Application-Grant Lag	1.9222	1.9413	0.492	1.9222	1.9381	-0.42
Number of citations 1999	10.023	9.1576	2.52**	10.023	10.344	-0.48
Γ_{1999}	.53217	.49939	9.7 ***	.53217	.53531	-0.20
Net Income	1404.2	676.5	9.74***	1404.2	1314.5	0.68
R&D investment	1702.7	1577.5	2.24**	1702.7	1768.3	-0.75
Sales	35140	26181	8.92***	35140	35669	-0.30
U.S. Corporation	.85111	.74143	7.48***	.85111	.852	-0.05
HJT subcatt2	.02333	.01278	2.79***	.02333	.01978	0.52
HJT subcatt3	.00667	.00351	1.58	0.114	.00711	-0.11
HJT subcatt4	.00111	.00803	-2.32 **	.00111	.00178	-0.37
HJT subcatt5	.00444	.02053	-3.40***	.00444	.004	0.15
HJT subcatt6	.08556	.05796	3.51***	.08556	.07985	0.44
HJT subcatt7	.14111	.13395	0.63	.14111	.14467	-0.22
HJT subcatt8	.12	.12053	-0.05	.12	.12222	-0.14
HJT subcatt9	.02222	.0547	-4.27***	.02222	.02244	-0.03
HJT subcatt10	.05556	.08727	-3.36***	.05556	.05578	-0.02
HJT subcatt11	.02	.02096	-0.20	.02	.02311	-0.45
HJT subcatt12	.00667	.03262	-4.38***	.00667	.00578	0.24
HJT subcatt13	.02444	.04984	-3.49***	.02444	.02822	-0.50
HJT subcatt16	.01778	.03246	-2.48***	.01778	.01511	0.44
HJT subcatt17	.01778	.01395	0.97	.01778	.01578	0.33
HJT subcatt18	.05778	.02143	7.41***	.05778	.06281	-0.45
HJT subcatt19	.03222	.01169	5.63***	.03222	.03444	-0.26
HJT subcatt20	.07111	.0463	3.51***	.07111	.06711	0.33
HJT subcatt21	.08889	.09245	-0.37	.08889	.07911	0.75
HJT subcatt22	.01444	.03383	-3.21***	.01444	.01578	-0.23
HJT subcatt23	.01111	.01626	-1.22	.01111	.00933	0.37
HJT subcatt24	.02111	.01317	2.07**	.02111	.02222	-0.16
HJT subcatt25	.03667	.01728	4.40***	.03667	.042	-0.58
HJT subcatt26	.04333	.01742	5.85***	.04333	.04644	-0.32
HJT subcatt27	.00444	.01439	-2.50**	.00444	.006	-0.46
HJT subcatt28	.00889	.01742	-1.95*	.00889	.00867	0.05
HJT subcatt32	.00667	.00929	-0.82	.00667	.00778	-0.28
HJT subcatt34	.01	.00228	4.72***	.01	.00956	0.10
HJT subcatt35	.00556	.00189	2.49**	.00556	.00467	0.26
HJT subcatt37	.04111	.03582	0.85	.04111	.03844	0.29
N	70539	903		4232	900	

Table 9: Balance

Sample	MeanBias	MedBias
Unmatched	9.6	7.9
Matched	1.9	1.5

Table 10: Results of the CDiD estimator

	GENERALITY			
	Coeff.	se	t	p
Treatment-Procurement	.0359102***	.0110172	3.259462	.0011235
Time	.0775832***	.0049311	15.7335	1.65e-54
Number of cites	-.0003439**	.0001357	-2.534912	.011277
_cons	.5383647***	.002176	247.4087	0
N	5135			

Table 11: Robustness check: Results of the CDiD estimator for the Net_Generality outcome

	Net_GENERALITY			
	Coeff.	se	t	p
Treatment-Procurement	.0585177***	.0113278	5.165857	2.48e-07
Time	.0790914***	.0049485	15.98283	3.74e-56
Number of cites	-.0004303***	.0001377	-3.125729	.0017836
_cons	.5392586***	.0021894	246.3093	0
N	5135			

Table 12: Results of the CDiD estimator for different measures of generality

	GENERALITY_IPC			
	Coeff.	se	t	p
Treatment-Procurement	.0574149***	.0112526	5.102369	3.48e-07
Time	.0889822***	.0052292	17.01654	3.29e-63
Number of cites	-.0006919***	.0001604	-4.313417	.0000164
_cons	.4890606***	.002381	205.4015	0

	GENERALITY_HJT			
	b	se	t	p
Treatment-Procurement	.0498474***	.0105924	4.705959	2.59e-06
Time	.0482736***	.0046793	10.3165	1.04e-24
Number of cites	-.0003184**	.0001384	-2.300911	.0214369
_cons	.4418676***	.0021578	204.7775	0
N	5135			

Table 13: Robustness check: Results of the CDiD estimator for the fake-treatment group

	GENERALITY			
	Coeff.	se	t	p
Fake-Treatment	.005146	.0102683	.5011556	.6162832
Time	.0766444***	.0050333	15.22744	3.04e-51
Number of cites	-.0002901**	.0001443	-2.009723	.0445129
_cons	.5367823***	.0022836	235.0647	0
N	5118			

Table 14: Robustness check: Results of the CDiD estimator for patents owned by different kinds of assignees

	GENERALITY			
	Coeff.	se	t	p
Treatment-Procurement	.0527449***	.0083128	6.344996	2.34e-10
Time	.0900577***	.0045028	20.00055	4.78e-87
Number of cites	-.0005597***	.0001672	-3.348333	.0008164
_cons	.5277459***	.0026812	196.8294	0
N	8722			