More of the same or something different? Technological originality and novelty in public procurement-related patents

Emilio Raiteri[a]

[a] *College of Management of Technology, Ecole Politechnique Fédérale de Lausanne*

Abstract

During the last decade demand-side innovation policies have received a renewed interest and innovative public procurement has been increasingly considered as a *de facto* technology policy both from researchers and policy makers. However, recent studies say little about the kind of innovations public procurement might be able to induce and about their technological impact. This gap in the literature is surprising if we relate the current debate on procurement with the economic-historical analyses illustrating the contribution of the government in spurring major technological breakthroughs in the US. This work hypothesizes that the innovative output induced by a public procurement contract is more exploratory and novel in nature compared to the counterfactual situation in which it was achieved in the absence of public demand. To empirically test this hypothesis I frame the problem in a quasi-experimental setting, in which patents are the units of analysis. Treated patents are the output of a US federal procurement contract. The control group is constructed through matching methods based on patent characteristics. I then compare treated and control group to check for differences on two kind of outcome variables: i) the originality index that builds on patents' backward citations; ii) a set of novelty-related measure based on the primary and secondary technological classification assigned to each patent. Results suggest that innovative public procurement produces innovations that are peculiar objects in the technology space, building on broader technological roots and embodying more novel and wider combination of technological capabilities.

Keywords: *Technological Originality, Novelty, Public Procurement, Patent Data, Technology codes*

1. Introduction

During the last decade demand-side innovation policies have received a renewed and growing interest. In particular, innovative public procurement has been increasingly considered as a *de facto* technology policy both from researchers and policy makers, especially at the supranational level (EU, 2010; OECD, 2013; Warwick and Nolan, 1994). The rising attention on the topic also prompted several empirical studies that began to evaluate the effect of innovative public procurement upon firms' innovative behavior, either measured in terms of input (i.e. R&D investment) or output (i.e. patents, share of turnover generated by innovated products).

However, these recent studies say very little about the kind of innovations public procurement might be able to induce and about their technological impact. This gap in the literature is quite surprising if we relate the current debate on public procurement as an industrial policy with the economic-historical analyses illustrating the government role in spurring major technological breakthroughs in the United States throughout the 20$^{\text{th}}$ century. Raiteri (2014) tries to fill this gap and to provide empirical evidence of the existence of a positive relation between procurement contracts and the rise in pervasiveness of given technologies through a patent data analysis. In particular he finds that receiving a citation from a patent related to an innovative public procurement contract significantly raises the generality of the cited patent. That paper, grounding on the extant literature, suggests, but do not investigate in depth, two potentially complementary explanations for this result: size and variety.

On the one side, public procurement might enlarge existing markets enough to induce additional improvements in the upstream technologies that eventually lead to larger adoption and increase in pervasiveness. On the other, it might stimulate a different kind of technological change, sparking the exploration of farther portion of the technology space and increasing the diversity of application of given technologies. The present work aims at evaluating the accuracy of the latter interpretation. In particular, I will hypothesize that the innovative output induced by a public procurement contract will be more exploratory and novel in nature compared to the counterfactual situation in which it was achieved in the absence of public demand.

To empirically test this hypothesis I frame the problem in a quasi-experimental setting, in which patents are the units of analysis. A patent is considered as treated if is the output of a procurement contract between a U.S. federal agency and a private enterprises based in the U.S.. The control group will be constructed through matching methods that exploit patents' observable characteristics to find suitable non-treated unit to proxy for the counterfactual situation in which treated units did not receive the treatment. I then compare treated and control group to check for differences on two kind of outcome variables. In the first place I follow the literature that tries to assess the basicness of university patents and I hence make use of the originality index first introduced by Trajtenberg et al. (1997), that builds on patents' backward citations. Given the limitations of this index I then also construct a set of novelty-related measure based on the primary and secondary technological classification assigned to each patent, as recently proposed by Strumsky et al. (2012), Akcigit et al. (2013) and Youn et al. (2014).

To construct the quasi-experiment I create an original dataset combining information from three sources. The main one is the *NBER patent data project*, a dataset providing abundant information about each patent granted by the USPTO between 1976 and 2006, together with the related citation data. The second one is the *USPTO Patent Full-text and Image Database*, which offers the full searchable text of every patent granted from 1976 onwards and that I will use to identify treated patents. Finally I use the *U.S. Patent Grant Master Classification File*, a file providing detailed classification information on all the patents issued by USPTO in the last 200 years.

In the next section I describe the literature that provide the motivation for this work and present the rationale of its main ideas. In section 3 I discuss how and why patent data might be used in this context and I present the formal hypotheses. Section 4 describes the data and the method used in the empirical analysis, while section 5 presents results and robustness checks. Conclusions follow.

2. Theoretical framework

2.1. Innovative public procurement as a 'de facto' technology policy

Demand has been long acknowledged as a major source of technological change (Schmookler, 1962; Kaldor, 1966). Despite the slowdown in the study of the link between demand and innovation occurred in the 80's, the demand side approach was never abandoned and has slowly but constantly regained attention (Von Hippel, 1988; Rogers, 1995; Malerba et al., 2007; Fontana and Guerzoni, 2008; Guerzoni, 2010). In this context, it has been increasingly advocated that also states, as particular kind of consumers, may be crucial for fostering innovation through the demand-side. The purchasing activities carried out by public agencies represents on average about 13 % of the GDP for OECD member countries (OECD, 2013; Warwick and Nolan, 1994) and, given its size, its technological composition is able to significantly affect the innovative behavior of the private sector. On this ground, *innovative public procurement* [1], has been increasingly acknowledged as a 'de facto' technology

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

policy (Cozzi and Impullitti, 2010), that may have similar or even greater effects on firms' innovative behavior compared to other more established supply-side policies such as R&D subsidies and tax credits (Geroski, 1990; Guerzoni and Raiteri, 2013).

According to Edquist and Hommen (2000b) innovative public procurement occurs 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 through additional or new technological development work, carried out in order to satisfy the demands of the buyer'.This form of purchasing is usually opposed to 'regular public procurement' which instead happens when the state buys simple products such as pen and papers, where no R&D is involved. As discussed in Raiteri (2014), several works criticized this definitions for being too restrictive (Uyarra and Flanagan, 2010; Rolfstam, 2012) and proposed broader interpretations that include also the pre-commercial stages in the process of product development, such as exploratory research services and prototyping (Edler and Georghiou, 2007). In this work (as in Guerzoni and Raiteri (2013) and Raiteri (2014)), following Rolfstam (2012), innovative public procurement is defined as all the purchasing activities carried out by public agencies that may lead to, or promote, innovation of some kind. This definition hence also includes what is usually described as pre-commercial procurement, usually a R&D service contract that may involve from exploratory research up to prototyping, as far as it produces a tangible innovative output[2].

While the rising attention on the topic stimulated the debate about the exact definition of innovative public procurement, most of the theoretical works (Geroski, 1990; Dalpé, 1994; Edquist and Hommen, 2000a; Edler and Georghiou, 2007) that speculate on the impact of public purchasing upon innovative input and output agree on the rational for considering procurement as a powerful innovation policy tool. Even if the first aim of public demand is not to stimulate innovation but rather to satisfy specific public needs (Edquist and Hommen, 2000a), public agencies, acting as qualified technology procurer, can in fact enlarge, or create, market for specific good and services ensuring the critical mass that may lead to larger R&D investment and to network effects that will spur the realization, adoption and diffusion of innovations. Moreover, since in fields such as defense cost considerations are often less important than performance, public demand can also stimulate innovation by providing firms with the opportunity 'to experiment with different product variants free from short-term commercial pressures'(Geroski, 1990).

In recent times policy-makers started to receive the recommendations about the relevance of demand-side innovation policy provided in these theoretical contributions. In particular, innovative public procurement is acknowledged as one of the main innovation policy instruments, together with support for R&D and capital market interventions, by several reports and policy initiatives by the European Commission (EU, 2010) and the OECD (OECD, 2013, 2014). While only a few countries, as for example Finland, have historically given high priority to innovation-oriented demand policies, innovative public procurement is currently receiving growing attention also at the national level, especially in emerging economies such China and Eastern Europe countries (Uyarra, 2013).

Despite the theoretical and political attention, most of the evidence on the actual effect of procurement policy is based on case studies(Edquist and Hommen, 2000a; Rolfstam, 2009; Uyarra and Flanagan, 2010; Flanagan et al., 2011; Brammer and Walker, 2011). However a few, mostly very recent, empirical studies have confirmed the theoretical hypotheses on the positive effect of innovative procurement on firms' innovative behaviour both in terms of input and output of the innovative process. For instance, Slavtchev and Wiederhold (2011), Draca (2012),and Guerzoni and Raiteri (2013) corroborate the results of an early work byLichtenberg (1988), providing econometric evidence of a reinforcing effect of public procurement on private investment in R&D. In the same way Aschhoff and Sofka (2009) and Draca (2012) find evidence of a positive effect of procurement in terms of innovation

[2]In the context of the present paper the filing of a patent document from the state contractor will proxy the innovative output of the procurement contract.

output, either measured through the share of turnover coming from the sale of innovated product and services of in term of patents.

Even though the role of public innovative procurement as an industrial policy is becoming more and more central in the debate among both the policy-makers and innovation scholars, the recent literature, regardless of its theoretical or empirical nature, say very little about the kind of innovation and technological change that innovative public procurement is able to induce.

While some empirical studies evaluate whether public procurement is more likely to lead to product, service or process innovations (Georghiou et al., 2013), very limited work (Uyarra and Flanagan, 2010) has been done to develop a theory that links the specific characteristics of innovative public procurement as innovation policy with the technological impact of the innovations it brings about. As discussed in Raiteri (2014), this gap is quite striking if we relate the current debate on the role of public procurement in spurring innovation with the economic and historical contributions that analyzed the development of major innovations and the growth of key industries in the United States during the 20^{th} century. A substantial body of studies (Nelson, 1982; Levin, 1982; Mowery and Rosenberg, 1982; Katz and Phillips, 1982; Flamm, 1987; Langlois and Steinmueller, 1999; Mowery, 2011, 2012) highlights in fact that the sheer size of federal procurement for purposes of national defense and spatial exploration had an essential impact in fostering radical technological breakthrough in industries that were key to the U.S. economic growth during the second half of the last century, such as the semiconductors, the computer, and the aviation industries. The work of Ruttan (2006) proposes an even more clearcut interpretation, arguing that defense related procurement has been the single most important factor for the development of every general-purpose technology (GPT) ever created in the United States, from mass production system to the internet and satellite communications, passing through the computer, semiconductors and nuclear power.

Despite the abundant qualitative evidence provided in these economic-historical analyses, also in this case very limited work furnished a structured theory or econometric evidence about the existence of a link between federal procurement and the development of major and pervasive innovations such as GPTs. Raiteri (2014) tried to fill this gap. The paper bases on the idea that the degree of pervasiveness of a given technology is not a steady characteristic but a dynamic attribute that can be evolved and cultivated over time by inducing investments in R&D and innovation in the sector that apply that technology. On this ground, he hypothesizes that innovative public procurement might be a powerful tool to prompt those innovation complementarities in the application sectors able to trigger the virtuous cycle that could eventually lead to an increase in generality of the upstream technology. Public procurement is hence thought as an enabler that can support promising new technologies early in their life-cycle, up to the point in which their efficiency improves so much that also a more risk averse private sector adopt them. Raiteri (2014) uses USPTO patent data, and in particular patent citations, to corroborate this hypothesis and finds evidence that receiving a citation from a patent related to United States' federal procurement rises the degree of generality of the cited patent with respect to the counterfactual situation in which such a citation did not arrive[3].

However, while Raiteri (2014) provides evidence of a solid relation between public procurement and the rise in pervasiveness of the technology on which federal procurement builds, he does not put forward that public procurement alone may generate major innovations but only that the innovative output of federal procurement contracts triggers a self-reinforcing process that lead to an increase in generality. The paper hence leaves some fundamental questions unanswered. In particular, it does not provide a conclusive explanation of why procurement output is able to set in motion the

[3]Raiteri (2014) designs a quasi experiment in which he compares the generality level (Generality Index) at two points in time in a group of treated and control patents. A patent is treated if it receives a citation from a patent related to public procurement. To mitigate endogeneity issues and to use non treated units to proxy for the counterfactual situation in which treated patent were not to receive the treatment, he implement the conditional difference-in-differences approach.

chain reaction that rise the generality of the technologies it piles on. On the one side it might only depend on a size effect. Public demand enlarges the size of an existent market enough to push firms to further invest in the development of a given technology. The consequential improvement, in terms of price or performance (or both), then spurs the diffusion and the adoption of that technology also in sectors that initially did not employ it because of the excessive cost or limited performance, increasing its overall pervasiveness. On the other, procurement could induce a radically different kind of technological change, setting in motion very different dynamics in terms of R&D investments and learning. Procurement may stimulate innovations that link previously not connected regions in the technology space, directly prompting more variegated application of a given technology with respect to what would occur in the absence of public demand. While in both cases the presence of public demand is pivotal to lead to the raise in pervasiveness of a specific technology, in the former case, it would not make any difference if the additional demand came from the private sector and not from the government. On the contrary, in the latter, while size would clearly still be a necessary condition, the sophistication and the peculiar preferences of the government as a consumer would be crucial to achieve variety (Malerba et al., 2007; Guerzoni, 2010) and the resulting increase in generality of a given technology. The rest of the present paper investigates the plausibility of the variety increasing effect by looking at the nature of the innovative output induced by federal procurement contracts in the United states.

2.2. Procurement innovative output: is only "more of the same" or is it something different?

As described in the previous section innovative public procurement might be considered as an effective technology policy tool, able to foster innovation and to increase the technological impact of specific technologies upon the whole economy. However, it is not clear which characteristics of the innovative output of the procurement process' is pivotal to achieve this result. Both the recent procurement literature and the economic-historical analysis devoted to the study of technical change in the U.S. agree on the important role played by the size of the state as a consumer in inducing innovation and major technological breakthrough. Geroski (1990) suggests the government may in fact enlarge initial market to the 'sufficient size to get the ball rolling'. As mentioned above, since state's demand could be large enough to stimulate further investment in R&D and technological development which in turn may accelerate widespread adoption and, in a virtuous cycle, further investment in R&D and technological advances.

A similar argument is borne out by the economic-historians. Mowery (2012) (among others) puts forward for example that the main reason for the successes of defense-related procurement projects in spawning radical innovations and civilian spin-off with respect to similar programs in other OECD countries such as UK and France, relies in the massive scale of the U.S. defense-related R&D and procurement programs. In order to properly understand the unambiguous impact of this scale-effect, it is enough to consider that, in the early years of development of the transistor and integrated circuits, government demand absorbed more than 50 per cent of the total industry output (Levin, 1982) and that the same holds true for the computer industry between 1945-1955 (Flamm, 1987). Moreover, as stressed by Mowery (2011), the size of defense-related procurement programs also abundantly increased competition in such emerging industries. The prospect of large contracts operated as a strong attractor for new firms to enter in the developing sectors and the tough competition further accelerated improvements in the price/performance ratio.

However, even if the scale of procurement programs was crucial for the deployment of several GPTs, another characteristic of the innovations induced by public procurement contract may have had crucial importance in achieving major technological breakthrough: their diversity. If we conceive the evolution of a specific technology as a sort of evolutionary tree, as originally suggested by David (1988) and Foray (1997), in which some branches are explored while others remain largely ignored, path-dependence and localized-learning processes may lead to the incomplete exploration of the variety distribution of that technology. Building on this idea, Cowan and Foray (1995), discussing the relevance of defense related R&D in an era of declining defense budget due to the end of the Cold war, hypothesize that the main contribution of military-sponsored research is to explore a different portion of the technological spectrum with respect to what civilian R&D would do. Private firms, in the absence of public intervention, tend to focus on immediately profitable innovations, avoiding distant search and experimentation. They will hence keep doing R&D for improving on already known technologies for which is possible to compute expected returns on the investment. In the evolutionary tree metaphor, private R&D will only explore a limited number of branches and due to localized technological change (Antonelli, 1995) and positive feedbacks, the discovery of new and distant ones would become more and more unlikely. To use the March (1991) categories, private firms will focus on the exploitation of old certainties and avoid ambiguous outcomes that result from exploration strategies. March underscores that:

> *The essence of the exploitation is the refinement and extension of existing competences, technologies, and paradigms. Its returns are positive, proximate, and predictable. The essence of exploration is experimentation with new alternatives. Its return are uncertain, distant, and often negative.*

On the contrary, as Geroski (1990) and Edler and Georghiou (2007) put forward, the primary interest of the state in its procurer activity is not the economic feasibility of a project but the satisfaction of specific social or political needs, such as national or energy security. Public agencies in procuring new technologies may hence act as experimental users (Malerba et al., 2007) with a strong preference for the future performance of a specific technology rather than for its present (lower) cost. Creating niche, or fringe, markets that are not well served by the existing technologies, public demand may stimulate innovation arising from more distant search in the technology space. Innovative public procurement, by providing safe, non-negative returns to exploratory R&D, would in fact bear part of the high uncertainties and ambiguities involved in the development of more original innovation that the private sector alone is not willing to face(Mazzucato, 2013)[4]. The innovative output of public procurement related R&D would then be more exploratory in nature and would increase more the variety of application of a new technology compared to what private R&D would do by itself.

Clearly, the size and the peculiar preferences of the state as a consumer should not be seen as substitutes in the process of increasing the diversity of the applications of a given technology. As proved in Guerzoni (2010), a minimum market size should be complemented with a high degree of consumers' sophistication and a good knowledge of the consumers' needs to achieve variety. The main aim of this paper is therefore not to evaluate the size effect hypothesis against the variety one, but rather to investigate whether the peculiar preferences of the government in its procurer activity are actually able to generate an unconventional kind of innovative output. I will hence try to corroborate the hypothesis that innovative public procurement produces not only 'more of the same' kind of technological change that the private sector would bring on, but generates a technological output systematically different in nature with respect to the one that private firms would realize in the absence of public demand. In particular, I will put forward that procurement related research, through a wider exploration of the technology space, will result in more basic and novel outcomes than corporate R&D.

In order to define the basicness and novelty I build on the long acknowledged view of invention as a cumulative process of recombination of the extant technical and scientific knowledge (Schumpeter, 1934; Nelson, 1982; Arthur, 2009; Weitzman, 1998; Fleming, 2001). If every innovation piles on pieces of previously existing knowledge and constitutes a new brick on which future innovations can build

[4]As described in Raiteri (2014), the state can directly take charge of the cost-related uncertainties. For instance, in the U.S. 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 reimburses the contractor the realized cost and pays an additional fee.

on, it is then possible to assess the degree of originality and novelty of a specific innovation by looking at which kind of knowledge and capabilities it mixes to be achieved. An innovation will be hence considered as more original when it builds on many different pieces of previously existent knowledge and when it combines in new and peculiar ways blocks of knowledge that are more distant in terms of technology space,. As a proxy for the output of the innovation process I will make use of patent data. In the next section I will describe how patent data has been and could be used in this context and how I can exploit the information they include to formally state my hypothesis.

3. Measuring originality and novelty through patent data

Patent data have long been praised as a fundamental source of information for researcher studying the economics of technological change(Schmookler, 1962; Scherer, 1965; Pakes and Griliches, 1980; Griliches et al., 1988).The main reason for their relevance is that patents can be considered as a direct output of the inventive process, and more specifically they represents outcomes that are expected to have an economic and commercial impact. Patents are in fact temporary monopolies granted to inventors/assignees for the commercial use of a new product or process in exchange of full disclosure. They have legal value and must include detailed information about the innovation itself, the inventors, the assignees, and about the citations to previous patents on which the invention builds on. Moreover patents have been around for quite a long time[5] and patent data are hence available in large numbers and for very long time series. For instance, in the U.S. patents have been granted continuously since the 1790's (Hall et al., 2001). However, patents are not a perfect measure of the inventive activity output since, as already stressed by Kuznets (1962), not every invention is patented. On the one side, an innovation may not be patentable because it does not meet the criteria of novelty, non-triviality and applicability set by the different patent laws. On the other, when the inventor has to make the decision whether to patent or not he may act strategically, deciding to keep the secrecy of his invention.

Even if patented inventions represent only a portion of the whole innovation process' outcome, patent data are probably still the best and richest source of information about innovative output. For this reason information coming from patent data has been used in many different ways to build diverse measure of technological change. Especially relevant for the present paper is the strand of literature that uses patent characteristics to assess the degree of basicness or originality in patents assigned to university or to academic staff (Trajtenberg et al., 1997; Henderson et al., 1998; Rosell and Agrawal, 2009; Thursby et al., 2009; Czarnitzki et al., 2009; Guerzoni et al., 2014). Most of the works in this field build upon an idea very related to the one exposed in the previous section: research performed by universities should be more basic, driven by scientific curiosity, and more exploratory in nature than corporate R&D. Consequentially, patents owned by universities should embody distinct features with respect to the ones assigned to private firms. In order to test this hypothesis Trajtenberg et al. (1997) and Henderson et al. (1998) developed new patent-specific measures computed on the basis of backward and forward citations included in the patent documents. While most of these studies find a significant (but declining over time) difference in the basicness measures between university and corporate patents , Guerzoni et al. (2014) stress that also the nature of the funding of a specific research inside the university affects the originality of a patent. Patents invented by academic inventors but partly funded by corporate money result to be less original than wholly university-sponsored patents. The latter point is particularly relevant for the issue raised here since I put forward that the difference in patents owned by similar institutions (i.e. private corporations) depends on whether they were induced by a public procurement contract or were the results of a completely private initiative. In this paper I hence follow the works that tried to evaluate the basicness of the output of university's

[5] The first patent in the world history was granted in Florence to Filippo Brunelleschi for the 'Badalone', a flat-keeled boat with paddle wheels for transporting heavy goods along rivers. The first patent law was instead implemented in Venice back in 1474.

research. I use patents as a proxy for procurement induced innovation output and specifically patent backward citations to explore the originality dimension.

3.1. Backward citations and originality

Patent citations have an important legal and economic meaning: if a patent cites another one, it implies that the latter constitutes a piece of knowledge upon which the former builds and over which it cannot have any claim (Hall et al., 2006). They are therefore compulsory and determine the scope of the property right awarded to a patent[6]. The idea that citations represent an extremely valuable piece of information to assess the quality of a patent dates to the contributions of Trajtenberg (1990) and Jaffe et al. (1993). It was only Trajtenberg et al. (1997) though that acknowledged that citations can say also something about the degree of basicness of a patent. Trajtenberg et al. (1997) noted that, since patent offices assign patents to specific technology class[7], citations between patents could be seen as linkages between different regions of the technology space and may therefore provide information about the originality a patent. According to Trajtenberg et al. (1997), if a patent cites previous patents belonging to a wide range of different technology classes, the broader will be the technological roots of the research that produced it, and the more original the patent. If instead all the citations made by a patent go to patents belonging to a single technology class, it means that it build upon a very specific set of technologies and is then less original. On the ground of this idea, they constructed the *Originality Index*, which is computed on the basis of the degree of concentration of backward cites across patent classes. While I will describe the index in details in section 4.1.3, it is important to note here that patent data and in particular backward citations allow me to compute a measure for the degree of originality of a patent and hence to formalize the idea presented in the previous section in the following hypothesis:

> *Hypothesis 1: A patent induced by a public procurement contract will have a higher originality index compared to the counterfactual situation in which it was the output of a completely private research project*

Even if the originality index is widely used in the literature and provide a concise measure for patents' basicness, it also presents several drawbacks. In the first place, since it is based on citations made by a patent, by definition it is clearly not computable for patents that do not include any citation to relevant prior art[8]. This might constitute a major issue in the context of the present work since it would force me to evaluate the originality of a group of patents by systematically excluding those patents that are potentially more novel. The fact that neither the inventor, nor the patent examiner are able to find any relevant prior art, may indeed means that a specific patent is the outcome of a very exploratory research process that reached the frontier of the known technology space. Secondly, the originality index is based on the concentration of citation across patent classes but it does not take into account the technological distance between patent classes, considering classes as equidistant one from the other (Hall and Trajtenberg, 2004). This is clearly an oversimplification since, for instance, information technologies classes are clearly much closer to communication technologies classes than to agriculture-related technologies. Also this issue could obviously represent a serious weakness for

[6] In the U.S. the applicant has the formal duty to disclose any knowledge of the prior art of which he is aware, and then a patent examiner verifies whether all the relevant prior art have been included and, if needed, adds missing citations. At the European Patent Office (EPO) all citations are instead recorded by the patent examiners who add the minimum number of citations to cover prior art .

[7] The USPTO uses about 480 major technology classes, of which only about 450 could be used as primary and unique, technology class.

[8] As we will see in section 4.1.3, when we correct for the index for the bias due to the low number of citations made, as shown by Hall (2005), we can compute the originality index only for patents that include at least 2 citations to prior art.

the correct evaluation of the originality of a patent, because patents citing very distant classes in the technology space may have the very same originality index as patent citing distinct classes but very close in terms of technological distance. Finally, the originality measure, as explicitly pointed out by Trajtenberg et al. (1997), is a backward-looking only measure. This means that it provides little information about the new connections in the technology space that a specific patent is bringing on. For example two patent documents that cites the very same patents belonging to an ICT-related technology class as prior art, will have the same degree of originality according to the originality index, even if one of the focal patent belongs to a class that was never linked to an ICT class before and the other exactly belongs to the same ICT-related class. Since the former case would clearly signal a broader and less standard search in the technology space than the latter, it could again represent a problem for the evaluation of the actual originality of a patent.

3.2. Technological classification, technology codes, and the novelty dimension

In order to overcome the potential weakness of the originality index, I follow a very recent body of works (Strumsky et al., 2011, 2012; Akcigit et al., 2013; Youn et al., 2014) that suggest to exploit patent technological classifications to better analyze the nature of technological change. Also this strand of literature clearly sees technological change as a process of recombination of new and existing capabilities and builds upon the seminal contribution of Schumpeter (1934) and the more recent ones by Nelson (1982), Weitzman (1998), andArthur (2009). Strumsky et al. (2012) first put forward that the use of patent citations to assess the degree of novelty embedded in a patented invention might be complemented by the analysis of additional patent metrics based on a finer-grained technological classification. In particular they suggest that technology codes assigned to a patent could be extremely fit to characterize the recombination process that lead to an innovation and to assess the degree of novelty embedded in it. Before going any farther it is hence worth briefly considering how technology codes are defined and assigned to a patent.

As we already mentioned, in order to make the search for relevant prior art easier for patent examiners, the USPTO assigns to patent one or more technological classifications. These classifications are numerical codes composed of two parts: a technology, or patent, class and a subclass. Technology classes are broader categories, while subclasses are detailed subsets of technologies within each class. To appreciate the magnitude of the difference in precision between the two kind of categories it is enough to report that, while the USPTO currently uses 480 major patent classes, there are more than 150,000 active subclasses (Strumsky et al., 2012). The technology code results from the combination of the patent class and the subclass. They are assigned to patent on the basis of the claims included in the patent document, because the claims state in technical and precise terms the subject matter of an invention and specify what is novel and patentable about it. Each claim can receive a different technology code and a single claim may also receive more than one technology code. Multiple technology codes may hence be assigned to the same patent and, while one of them should be chosen as the primary (or original) classification, there is no limit as to how many codes may be assigned to a patent[9] (Strumsky et al., 2012). As described in the next session, in the sample I use for this analysis more than 90% of the patents embodies at least 2 technology codes, whereas the average number of codes per patent is 4.4. Patents can be seen as a combination of different technology codes and, since these codes describe the features of existent technologies, a patented invention represents an 'instantiated combinations of technologies from the defined set of all possible combinations of technologies' (Strumsky et al., 2012). Given that when the USPTO adds a new code to identify the appearance of a new technology, a retroactive reclassification of all previous patents that may incorporate that technology takes place[10], USPTO technology codes represent a set of consistent definition of technological

[9]For a more detailed description of the process through which technology codes are assigned to claims and then to patents see Strumsky et al. (2011) and Strumsky et al. (2012).

[10]on the reclassification issue

capabilities spanning more than 200 years of inventive activity (Youn et al., 2014)[11]. Exploiting this features of technology codes, Strumsky et al. (2011), Akcigit et al. (2013) , and Youn et al. (2014) put forward that combinations of codes can tell us something about the degree of novelty of a patent. Looking at new patented inventions in terms of codes combinations, it is possible in fact to categorize innovations in three main groups:

- *Novel technologies*: innovations that arise from novel combinations that involve a completely new technology code and hence new technological capabilities.
- *Novel combinations*: innovations that embodies new combinations of already existing technology codes, meaning that a specific combination of technology codes was never observed in the patent office's history, even if the single technology codes are not new.
- *Refinements*: innovations that incorporate old combinations of technology codes, which means that the very same combination of codes was already used to describe the technological content of an earlier patent.

Although the approaches used in the afore mentioned studies to put a patent in one of these three categories are slightly different, in all of them the data clearly shows that the process of inventions from the second half of the 20^{th} century has been almost completely driven by novel combinations (70 to 80% of all patents, depending on the application year[12]) and refinements (20-30%), with novel technology accounting for less than 0,5% of the total patent population since the 1970's. A new patent can then almost only stem from a novel combination or from a refinement.
Moreover Youn et al. (2014) propose a further differentiations between *narrow* and *broad* combinations, based on the technological distance within patent classes. Since patent classes are major categories and subclasses are much more detailed subsets, a combination of technology codes is defined as *broad* if the technology codes used to classify a patented invention belong to different technology classes (i.e. they are more distant in term of technology space), it is instead defined as *narrow* if all the codes assigned to a patent belong to the same patent class.

The notions of *refinement* versus *new combination* and of *broad* versus *narrow* combination provide then an additional way to assess the novelty of a patent. In section 4.1.3 I will develop a set of novelty measures that will build upon these concepts and in particular on the idea that more exploratory research will more likely lead to innovations that embody novel rather than refinements on old combinations, and broad rather than narrow combinations. In the framework of the present paper these measures therefore allows me to formulate the following additional hypotheses that complement the one on patent originality:

> *Hypothesis 2a: Patents induced by a public procurement contracts will embody a larger share of novel combinations compared to the counterfactual situation in which they were the results of private research projects*
>
> *Hypothesis 2b: Patents induced by a public procurement contract will embody a larger share of broad combinations compared to the counterfactual situation in which they were the results of private research projects*

[11]The first patent registered at the USPTO dates back to year 1790.
[12]discrepancies between youn and Akcigit

4. Data and Method

In order to test the hypotheses stated in the previous section I frame the problem in a quasi-experimental setting in which the single patent is the unit of analysis. I therefore compare the degree of originality and novelty between a group of treated and a group of control patents. A patent belong to treated group if it has been induced by a procurement contract between the federal government of the United States and a private corporation. The control group will be instead cautiously constructed in order to achieve the highest similarity with the treated group in terms of observed patent-specific characteristics.

4.1. Data and sample selection:

I created an original database exploiting information coming from three different sources. The first one is the *NBER patent database*[13]. It contains detailed information about 3,209,376 unique patent granted by the USPTO from 1976 to 2006 and specifically: 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 primary U.S. 3-digit patent class and subclass, the number of claims made by each patent, and backward and forward citations.

The second source is the *USPTO Full-text and Image Database*, which offers the full searchable text of every patent applied granted from 1976 onwards and allows to look for specific words and pieces of text within the different fields of the patent documents[14].

The last source is the *U.S. Patent Grant Master Classification File*, a file that contains the primary and all the secondary classification information on all the patents issued by USPTO since the 1790's[15]. Both for the primary and the secondary classifications, the data report the full technology code composed by the major patent class and the subclass.

As in Raiteri (2014), given the abundant information included in it, I use the NBER database as the main source of information about patents and, in particular, I exploit backward citations data to construct the first outcome variable, the originality index. The Master Classification File instead provides the data for the construction of the novelty-related outcome variables, based on the combination of technology codes. Finally, the USPTO full text database is used to identify the patents induced by public procurement contracts. The next section carefully describes how I use the three dataset to build the structure of the quasi-experiment.

4.1.1. Sample selection

The NBER patent database includes information about more than 3 millions patents. From this data set I draw only patents that are assigned to private corporations based in the United States and that were applied for between 1985 and 2004. The rationale for considering only patents applied for after January 1st 1985 is that, as explained in section 4.1.2, the law I exploit to identify patents induced by procurement contracts was introduced in 1983. Considering patents from 1985 onwards hence ensures that the new rule was fully in force and well understood by all actors during the period taken into account. I also exclude the last 2 years for which data are available, 2005 and 2006,

[13] The NBER patent dabase is the result of the effort of different researchers, Hall et al. (2001) developed the dataset and carefully describe the information included in its first version, which has been continuously updated since. Data are available at `www.sites.google.com/site/patentdataproject`.

[14] The full list of the fields and the online version of the dataset are available at `www.patft.uspto.gov`.

[15] Google and the USPTO have entered into an agreement to make some USPTO products available to the public at no charge. The Grant Master Classification File is among them and is available for download from Google USPTO Bulk Dowloads at `http://www.google.com/googlebooks/uspto-patents-class.html`.

Table 1: Descriptives statistics for the whole sample

	Mean	SE
Patent characteristics		
App-grant_lag	2.20	1.18
# claims	18.02	14.51
# codes	4.49	3.36
# cites_rec	8.69	16.11
# cites_made	12.12	20.44
Patent_stock	4855.38	9684.72
Technological category		
Chemicals	0.17	0.37
ICT	0.22	0.42
Drugs and Medicals	0.11	0.31
Electrics and Electronics	0.20	0.40
Mechanical	0.15	0.36
Others	0.15	0.36
Time windows		
1985-1989	0.18	0.39
1990-1994	0.21	0.41
1995-1999	0.33	0.47
2000-2004	0.28	0.45
N	1,039,052	

to reduce potential bias due to truncation. Since the average lag between the application and the granting of a patent is about 2 years, and because less original and novel patents need less time to be examined and then granted, refinements and less original patent might be overrepresented in the last years[16]. The reason for constraining the analysis to patents owned by private firms depends instead on the very nature of my research question. Since I am interested in understanding whether patent induced by federal procurement contract have a higher degree of originality and novelty with respect to the patents that the system would have produced in the absence of public demand, it would be harder to identify the effect of public procurement if we included in the analysis patents assigned to university or research institutes that a considerable amount of works already recognized as more original and basic in nature than corporate patents[17]. With these limitations the full sample I use in the analysis counts 1,039,052 patents. Given that the NBER dataset only provides information on the primary technology class and subclass (i.e. only the primary technology code), I then merge the data with the Master Classification File, to obtain all the secondary technology codes assigned to a patent. After the merging I then have, for each patent in the sample, detailed information about: the application year, the grant year, the primary and secondary classification information, the number of claims, the number of technology codes assigned to a patent, the number of citations made and received, the primary technology class of citations made and received. Even if I do not have any additional information about the assignees, I also compute the patent stock as the total number of patents assigned to the same private corporation at the end of the period, to have at least a distant proxy for the size of the firm.

Table 1 reports descriptive statistics for the whole sample. As the table shows, most of the patents were applied for in the last ten years of the considered period, a fact that is consistent with the patent inflation phenomenon. Most of the patents belong to the ICT and the Electric and Electronics technological categories[18]. Each patent made on average 12.1 backward citations and has been encoded

[16]Removing only the last 2 years could be not sufficient to completely eliminate the truncation bias. To be sure that the truncation issue is not bringing about severe over or underestimation, in section 5.1, I will implement a robustness check stratifying data by 4 time windows of five years each.

[17]Further research could be done in evaluating the effect of procurement on university assigned patents

[18]Technological categories are defined according to Hall et al. (2001), and therefore correspond to the HJT categories.

with 4.4 different technology codes.

4.1.2. Treatment variable

Once I have selected the sample I use in the analysis, I have to identify the patents that compose the treatment group. As described at the beginning of this section, a patent is put into the treatment group if it was induced by a procurement contract between a U.S. federal agency and a private corporation. As in Raiteri (2014), in order to pinpoint those specific patents, I exploit the U.S. Federal Acquisition Regulation (FAR) and the *USPTO Full-text and Image Database* described in the previous section.

The FAR regulates the acquisition process by which the government purchases goods and services and includes specific articles that govern intellectual property rights in federal government contract. Although initially approved in 1974, it was only with the 1980's the FAR developed a set of homogeneous rules on property right management across all federal agencies and contractors. The FAR was first updated in 1980 to follow the Bayh-Dole Act, requiring that each contractor may elect to retain title to any invention made, or first reduced to practice, in the performance of work under a Government contract[19](FAR 27.301) (Sharp, 2003; Bloch and Gray, 2012). While this update in the FAR was originally addressed only to small business and non-profit enterprises, the president Memorandum issued by Reagan in February 1983 extended its scope to large and for-profit firms (Sharp, 2003; Bloch and Gray, 2012) to increase contractors' willingness to collaborate with the federal government (Sharp, 2003). From then on the FAR required that any kind of contractor can decide to retain the title to an invention, but also that, if it chose to retain title, it must file a patent application, granting to the the government 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). In order to ensure the license to the government, the FAR obligates the contractor to include into the patent application document a special clause called government interest statement. In this statement the contractor declares that the invention was made with Government support, and that the government has certain rights in the invention (FAR 52.227-11).

As mentioned in section 4.1, the USPTO Full-text and Image Database allows to search for specific pieces of text in different fields of the patent document, among which *Government Interest Statement*. A typical goverment interest statement would report a sentence like the following:

> *This invention was made with Government support under Contract/Grant Number [x] awarded by the Department [y]. The United States Government has certain rights in the invention*

It is hence possible to identify the patents including the government interest statement and to disentangle between the patents that include the statement because were induced by a grant, and the ones that make explicit reference to a contract between an organization and a federal department or agency. I then define my treatment variable *Treatment_Procurement* as taking the value 1 if a patent includes the government interest statement and mentions the word 'contract' and not the word 'grant', and the value 0 otherwise[20]. In this way I identify the 9,524 patents, induced by a public procurement contract between a private corporation and the federal government, that compose my treatment group.

Table 2 displays the descriptive statistics for the treated patents. As the table shows, treated patents are more concentrated in the electrics and electronics and chemicals technological fields with respect to patents in the whole sample. They also make and receive less citations on average, but

[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 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) used a similar strategy to identify patent resulting from grants of the National Institute of Health (NIH).

Table 2: Descriptives statistics for patents in the treated group

	Mean	SE
Patent characteristics		
App-grant _lag	2.22	1.20
# claims	18.97	15.21
# codes	4.52	3.19
# cites_rec	7.81	12.63
# cites_made	9.97	12.69
Patent_stock	4672.01	8716.08
Technological category		
Chemicals	0.20	0.40
ICT	0.16	0.37
Drugs and Medicals	0.03	0.16
Electrics and Electronics	0.36	0.48
Mechanicals	0.17	0.38
Others	0.08	0.28
Time windows		
1985-1989	0.21	0.41
1990-1994	0.20	0.40
1995-1999	0.30	0.46
2000-2004	0.29	0.45
N	9524	

include more claims and technology codes and belongs to corporation that have on average smaller patent stock than in the full sample.

The declaration reported in the government interest statement also allows to disentangle between the department awarding the contract that led to a specific patent. The pie chart in figure 1 clearly shows that most of the patents induced by procurement contract were awarded by the Department of Defense (51.7 %), the Department of Energy (37.2 %), and, to a lesser extent, by the NASA (7.8%). Only 3.3 % of the patents in the treated group were induced by contracts with other departments.

Figure 1: Share of procurement induced patents by department

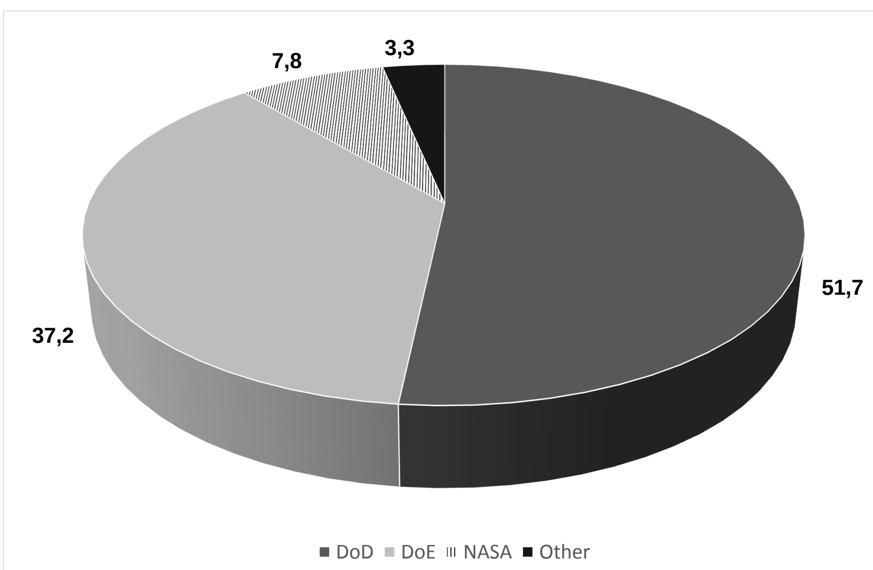

Figure 2 shows the distribution of treated patents across technology fields by the three department that awarded most of the contracts: Department of Defense (DoD), Department of Energy (DoE), and Nasa. As the descriptive statistics already showed, all the departments induced most of the patents in the Electric and Electronics field. However, it is also worth noting that, while patents induced by contracts awarded by the DoE are more concentrated in the Chemical technological category, the ones by the DoD have a concentration in the ICT sector.

Figure 2: Distribution of procurement induced patents by technological category and department

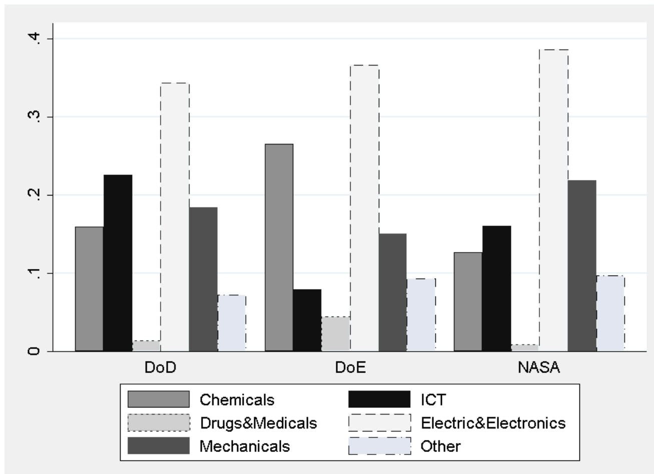

As in Raiteri (2014), a caveat should be made about this identification strategy. Even if the patents composing the treatment group are undoubtedly induced by public procurement contracts, they may not represent the whole universe of innovations brought about by procurement contracts with the U.S. federal government in the period taken into account. In special cases, federal agencies may in fact retain title on specific technologies or impose secrecy for purposes of national security. Moreover, a firm can try to keep an invention secret, and avoid filing a patent for it, or it could intentionally fail to report the government interest statement in the patent document. While it is clearly not possible to have a figure about how often the state imposes secrecy, two works evaluate compliance in reporting government involvement in inventive activities sponsored through federal research grants. A report by the Government Accountability Office conducted in 1999 (GAO, 1999) finds that for 10 to 20 per cent of the patents in their sample, grantees failed to add the government interest statement, even if it had to be included. More recently, Rai and Sampat (2012) analyze the reporting behavior in universities that received research grants from the NIH for biomedical research. Also in this case the authors find that a non-negligible share of patents that should include the statement do not report it: a share close to 40% during the 1980's, that reduced to 20-30% during the 1990's, and 10-20% in 2000's. If similar shares hold also for patents induced by contracts[21], looking at the half-full cup, I should be able to identify more than 70% of the total patents induced by public procurement in the 90's, and more than 80% in the early 2000's. Nevertheless, since I am clearly identifying patents induced by procurement contracts, potential failures in recognizing patents related to procurement contracts would threaten the results in two cases. The first one would occur if patents that should include the government interest statement and do not are systematically different in terms of originality with respect to the ones correctly identified as treated. The second would arise from using patents that should belong to the treated group (i.e. including the statement) to construct the control group. While it is hard to evaluate the relevance of the former case, in the latter, since procurement related patents, if anything, should be more original, the results would be biased downward and hence the existence of a difference

[21] Firms

between the two group would not be questioned. To take the first issue into account, since compliance appears to be rising over time, in section 5.1 I will implement a robustness check stratifying the data by 4 time windows of 5 years, to test whether the nonobservance of the FAR is biasing the result.

4.1.3. The outcome variables

As described section 3, I make use of two different kinds of outcome variables. The first one, the originality index, follows the literature about the basicness of university patents and builds upon backward citations. An additional set of novelty-related measures bases instead on technology codes. Even if originality and novelty may grasp slightly different latent characteristics of a patent, I use them as proxy measure to evaluate the breadth of the exploration process undertaken to reach a patentable invention, and hence to test the hypothesis presented in section 3.

The originality index

The originality index was first introduced by Trajtenberg et al. (1997) to measure the breadth of the technological roots of a patented invention. Hall (2005) then improved on the first version of the index, correcting it for the bias due to the count nature of citation data. I here use the corrected version of the index, which is defined as:

$$\Omega i = [N_i/(N_i - 1)][(1 - \sum_{j=1}^{J} N_{ij}^2/N_i)] \tag{1}$$

where *Nij* is the number of backward citations made by patent *i* to patents in technological class *j*, while *Ni* is instead the total number of backward references made by patent *i*. The summation term is the Herfindahl concentration index reporting the degree of concentration of backward citations across primary patent classes, measured at the 3-digit level (i.e. 470 classe). The whole second term of the equation reproduces the originality index as proposed by in the first place and the first term is the correction introduced byHall (2005). The *Originality index* is by definition bounded between 0 and 1: it gets closer to 1 as a patent make citations to patents belonging to a wide range of different patent classes, while it approximates 0 as its backward citations are concentrated in a few classes. Exploiting the information on patents primary classifications and on backward citations included in the NBER patent database, I therefore compute the *Originality index* for all the patents in my sample that made at least 2 references to relevant prior art.

The novelty measures

As discussed in section 3, the Originality index presents multiple drawbacks. In order to partially overcome some of them I develop a set of additional variables that try to grasp the novelty of a patent by looking at the combinations of technology codes used to describe the patents' technological content.

As highlighted by Strumsky et al. (2011), Akcigit et al. (2013), and Youn et al. (2014), at least form the 1960's, almost 99% of the patents stem from refinements of old combination of technology codes(25-30%) and novel combinations of existent technology codes (75-80%). Given this major differentiation, I hypothesized that patents induced by public procurement contracts should embody more novel combinations than refinements. In order to test this hypothesis and to determine whether a combination is indeed novel or not, I exploit the *Gran Master Patent File* presented in 4.1. For each combination of technology codes included in a patent I check if that specific combination has ever been used to describe any previously appeared patent in the history of the USPTO in terms of application date. If not, the combination is defined as novel, otherwise as a refinement. On this ground I then define the dichotomous variable *Novel* as taking the value 1 if a patent embody a novel combination of technology codes, and 0 if instead incorporates an old combination.

Youn et al. (2014) then proposed a further differentiation between broad and narrow combinations based on the technological distance within patent classes. Building on this intuition and on the notion of technological distance put forward by Trajtenberg et al. (1997), I develop three additional outcome variables that try to capture the breadth of novel combinations. Trajtenberg et al. (1997) developed a higher-level classification in which patent classes aggregated into 36 two-digit technological sub-categories (HJT-subcategories), and these in turn are further aggregated into 6 main categories (HJT-categories). They then suggested that technological distance between two patents is minimum when they belong to the same patent class, increases when two patents do not belong to the same class even in the higher level categories (HJT-subcategories), and is therefore maximum when two patents do not pertain to the same main HJT-category. Applying this notion of technological distance to the concept of broad combination proposed by Youn et al. (2014), I first define the binary variable *Low_Breadth*, that takes the value 1 if a novel combination mixes technology codes that belong at least to two different patent classes (3-digit level, 470 cells), and takes the value 0 if all the technology codes used to define a patent belong to the same patent class. I then define the variable *Med_Breadth*, taking the value 1 if a novel combination mixes technology codes that belong at least to two different HJT subcategories (2-digit, 37 cells), and 0 otherwise. Finally I create the variable *High_Breadth*, which takes the value 1 when a novel combination mixes technology codes that belong at least to two different HJT-categories (1-digit, 6 cells), and 0 otherwise. These four variables hence allow to test the hypothesis presented in section 3 and, in particular, to verify if patents induce by public procurement contracts embody more novel and broader combinations than control patents.

The novelty-related variables described above also provide us with a clear vision of a stylized fact highlighted in previous works. As pointed out by Strumsky et al. (2011) inventors in the last three decades seem to increasingly rely on exploitation of the known technology space, rather than exploring new portions or connecting distant regions of that space. Also Youn et al. (2014) put forward that is becoming harder and harder to sustain inventions on the basis of broader combinations. These trends are confirmed by the data in my sample. Figure 3 shows the share of patent involving both novel and broader combinations over time for the full-sample.

Figure 3: Share of patents with novel and broad combinations

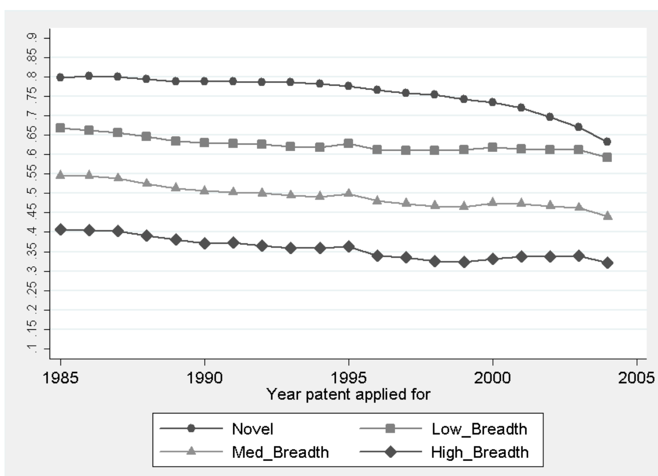

Even if truncation effects may have some role in shaping the trend in the last years of the time series, the share of patents with novel combinations constantly and remarkably reduced during the period taken into account. The same holds for breadth-related measures, the share of patents with broad combinations is consistently reducing over time, no matter the aggregation level I consider.

4.1.4. Descriptive statistics

Now that every component of the quasi-experiment has been defined, we can have a look at the descriptive statistics for the main patent characteristics that will be used in the analysis and for

Table 3: Means and t-test for mean differences in patent characteristics

Patent characteristics	**Not treated**	**Treated**	**Difference**
App-grant _lag	2.195	2.223	-.0287**
# claims	18.016	18.970	-.954***
# codes	4.492	4.522	-.029
# cites_rec	8.694	7.807	.887***
# cites_made	12.136	9.973	2.163***
Patent_stock	4857.07	4672.01	185.06*
Technological category	**Not treated**	**Treated**	**Difference**
Chemicals	.168	.200	-.032***
ICT	.223	.164	.058***
Drugs and Medicals	.1076	.025	.081***
Electrics and Electronics	.199	.355	-.156***
Mechanicals	.149	.170	-.021***
Others	.152	.082	.070***
N	1,029,528	9,524	

Outcome variables	**Not treated**	**Treated**	**Difference**
Originality	.538	.584	-.045***
N	938,073	8,648	
Novel	.759	.784	-.025***
N	1,029,528	9,524	
Low_breadth	.537	.542	-.005
Med_breadth	.419	.449	-.030***
High_breadth	.293	.342	-.048***
N	781,412	7,467	

the outcome variables. Table 3 displays the means and the result of a t-test of mean-differences by treatment status. As the table show, patents induced by public procurement contracts (i.e, the treated patents) present significant differences with respect to non-treated patents. They include on average almost 1 more claim, make .88 citations to prior art less and receive about 2 forward citations less than not treated patents. They also belong to assignees that have on average a larger patent stock at the end of the period. They are much more concentrated in the Electrics & Electronics technology field and less in ICT and Drugs & Medicals. In terms of outcome variable, treated patents appear to be on average more original and to embody a larger share of novel and broad recombinations, apart from the *Low_Breadth* variable, for which there is no statistical difference in the means for the two groups.

It should be noted here that, the five outcome variables are defined for 3 different samples of different size. While the *Novel* variable, building only on technology codes, is defined for the full sample, the *Originality Index* as mentioned in section 3, is determined only for patents that make at least two references to previous patented prior art and the sample used for analyzing this dimension counts 880,540 patents. The *Breadth* variables consider instead only patents embodying novel combinations, which, as table 3 shows, represents about 76% of the whole sample. The subsample used to estimate the effect on the *Breadth* outcomes, not considering refinements, counts to patents.

4.2. Empirical approach:

4.2.1. The evaluation problem

In section 3 I hypothesized that patent induced by procurement contract between a private corporation and the US Federal government, being the outcome of more exploratory research activities, should display higher lever of originality and embody larger share of novel patents with respect to what would have happened in the absence of public demand. In order to test these hypotheses I designed a quasi-experiment in which the patents induced by procurement contracts are considered as treated. Theoretically I am hence interested in estimating the effect of the treatment on the treated, the average treatment effect (ATT):

$$ATT = E[Y^T - Y^C|T] \tag{2}$$

Where Y^T represents our outcome variables (i.e. the originality index, or the novelty related measures) if treated, Y^C if not treated, and T specifies that all the patents considered here are the ones that belong to the treatment group. However, as usually in policy evaluation, since it is clearly impossible to observe the very same group of patent simultaneously as treated and as not treated, I am forced to use non-treated patents to proxy for the counterfactual situation. Formally:

$$ATT = E[Y^T|T - Y^C|C] \tag{3}$$

However, taking simple differences in the average between the treated and the non-treated group could lead to distorted results due to selection bias. As also the descriptive statistics reported in the previous section show, the treated and not treated group of patents are in fact very different from each other along several important dimensions that may affect the outcome. It is hence very likely that the treated group would behave in a different way with respect to the non-treated group even in the absence of the treatment (i.e. non zero selection bias):

$$E(Y^C|T) - E(Y^C|C) \neq 0 \tag{4}$$

In this case it is then not possible to correctly identify the average treatment effect since I cannot disentangle between the selection and the treatment effect. To consistently estimate the ATT I have hence to mitigate the selection problems. The literature dedicated to program evaluation suggested different methodologies to deal with the selection problems (Imbens and Wooldridge, 2009). In the particular context of the evaluation of the effect of a given policy on patent-specific characteristics, numerous works (Lanjouw and Schankerman, 2004; Czarnitzki et al., 2011; Feldman and Yoon, 2012; Fier and Pyka, 2012) carefully construct the control group including patents that are comparable to the treated ones along several observable dimensions. Raiteri (2014) follows this approach and adopts the matching method as a first step of a conditional difference-in-differences (CDiD) estimation. Since in the present setting the CDiD technique is not feasible because the originality and novelty-related variables are steady over time, I here exclusively rely on matching techniques, and then only on observable characteristics, to estimate the ATT. However, in order to reduce as much as possible the selection bias due to observables, I use a combination of two matching methodologies as in Lechner (2000): the propensity score and exact matching.

4.2.2. Matching methods

The basic aim of non-paramentric matching methods is to identify a group of non-treated units that are similar (ideally identical) to the treated ones in all the relevant covariates that affect treatment participation. When it is possible to find such a group, we can use it as a close substitute for the unobservable counterfactual situation in which the treated group is not receiving the treatment [22] and therefore to consistently estimate the ATT as:

$$ATT = E(Y^T|T, X) - E(Y^C|C, X) \tag{5}$$

[22] For a full review of matching techniques see Caliendo and Kopeinig (2008) and Stuart (2010).

where X is the vector of relevant covariates. However, matching methods allow the identification of the ATT only if a key assumptions is satisfied, the so called conditional independence assumption (CIA), or unconfoundedness (Rubin, 1973). It requires that the assignment to the treatment is independent of the outcomes, conditional on the set of observable covariates (X). Since for the CIA to be valid all the possible variables affecting the probability of being treated should be known, observed, and taken into account in the estimation, the vector of relevant covariates should be a high-dimension vector. In such a case, it is hence very hard to find matches with the same exact values for all the covariates, especially if those covariates are continuous in nature. In order to overcome this difficulty, Rosenbaum and Rubin (1983) introduced the propensity score matching method. This methodology allows, if the CIA holds, to condense the vector of relevant covariates into a single scalar index, the propensity score, which measures the probability of receiving the treatment given the observed covariates.

An additional assumption that should be satisfied is the common support condition, or overlap condition. It requires that the vector of relevant covariates cannot perfectly predict the assignment to the treatment or the control group, and ensures that there is sufficient overlap in the relevant characteristics of the treated and untreated units to find suitable matches. Formally:

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

If both conditions hold, the treated and the control group, once matched, should be on average observationally identical and it is then possible to retrieve the ATT, through the Propensity score matching estimator defined as follow:

$$Psm_{ATT} = E(Y^T|T) - E_{P(x)|T}[Y^C|C, P(X)] \tag{7}$$

Several works (Lechner, 2000; Puhani, 2002) highlighted that some variables may have a larger influence in determining treatment participation and propose a variant of the propensity score matching methodology in which exact matching on a few key covariates is combined with propensity score matching. The units of the treatment group are therefore paired in a sub-cell with units in the control group that have identical values for the key covariates and are then matched with those units in the sub-cell that are closest in terms of propensity score. This strategy, sometimes referred to as 'hybryd matching', ensures smaller selection bias since it avoid matches between units that present different characteristic in the basic covariates. Since, as I will explain in the next section, particular patent-specific characteristic may have a prominent impact in determining treatment participation, I will implement the latter variant of matching methods, combining exact and propensity score matching.

4.2.3. Exact matching and propensity score specification

A key step in using matching techniques to estimate the average treatment effect is to select the right variables to include in the vector of relevant covariates X. Caliendo and Kopeinig (2008) and Stuart (2010) recommend to include in the matching procedure all the variables known to affect both treatment assignment and the outcome variable, and to exclude variables that are affected by participation in the treatment. Given these requirements, covariates selection should base upon a good knowledge of the economic theory driving the process, previous research on the topic, and also information about the institutional setting that we wish to analyze (Caliendo and Kopeinig, 2008). As in Raiteri (2014), I hence pay attention at the peculiar nature of my treatment variable and I also largely follow the example of previous works that used treatment models for patent analysis (in particular to Lanjouw and Schankerman (2004), Feldman and Yoon (2012), Fier and Pyka (2012), and Czarnitzki et al. (2011)).

Because in this case I am mixing exact and propensity score matching, I first have to select the key variables used to estimate the propensity score and then the subset of covariates used in the exact matching procedure. Given that the propensity score represents the probability of receiving the treatment conditional on the relevant covariates, in this case, as in most quasi-experimental setting, it is clearly not observed and need to be estimated. As described above, the theory requires that, to consistently recover the propensity score, all the possible variables known to be related to both the treatment assignment and the outcome should be taken into account in the estimation process. To fulfill this requirement I hence estimate three different propensity scores, one for each subsample, through probit regressions of the treatment variable *Treatment_Procurement* on a set of relevant patent-specific characteristics. Consistently with the previous literature (Lanjouw and Schankerman, 2004; Fier and Pyka, 2012) I then include the following variables as covariates in the regressions: i)*Application Year Dummies.* Since it is very likely that the probability for a patent to be induced by a public procurement contract depends on the trends in public expenditure, the time dimension clearly affects the probability of being treated and I hence include 20 binaries reporting the patent's application year in the probit regression ; ii) *Technological subcategory dummies.* Public procurement contracts are awarded more often to firms that belong to specific sectors and patent more often in some technology class than in others. Consequentially, the technology class clearly affects the probability of a patent to be induced by a procurement contract and should be considered in the estimation. I then include 36 dichotomous variables accounting for the HJT technological subcategories (2-digit). iii)*Application-Grant Lag.* It is a variable that reports the number of years gone by between the filing date and the granting of the patent. The rational for considering this variable relies in the fact that the originality or novelty of the patent might be increasing with the length of the lag, because refinements and marginal improvements on existing technologies might be easier to evaluate by patent examiners; iv)*Number of claims.* As Hall et al. (2001) put forward, the number of claims in a patent can grasp information about the scope, or the width, of the monopoly power granted to a patent.Since the broader the scope of a patent, the more significant in terms of technology improvements an invention might be (Lanjouw and Schankerman, 1999), the patent scope may affect my outcome variables. The number of claims is then included in the probit regression as a proxy for it ; v) *Number of codes.* As for the claims, the number of technology codes used to describe a patent correlates with the scope of a patent. Moreover, the larger the number of technology codes in a patent, the higher the odds that it incorporates a novel combination of codes; vi)*Number of citations made.* Clearly, since the number of citations made by each patent depends on the width (existence) of the (any) relevant prior art, the number of backward references could be negatively correlated to the basicness of an invention and should be taken into account in the estimation; vii) *Number of citations received.* This variable reports the number of forward citations obtained by a patent at the end of the period for which data available are available (December 2006). Even if this variable is obviously not steady over time and might be affected by treatment participation, I include it in the probit regression since I consider it as proxy for patents' steady characteristics such as quality and economic value (Hall, 2005). As I hypothesized in section 3, procurement-related research activity might be less concerned with the economic value of its output with respect to business research. Our treatment variable might hence be negatively correlated with forward citations; viii) *Patent stock.* This variable reports the number of patents owned by each patent assignee in the sample (i.e. American private firms) at the end of the period in 2006. While also in this case the variable may be somehow affected by treatment participation, I use it to include at least a noisy measure for the size of the firms owning the patents, given that the data at our disposal do not comprehend any other information about the assignees. The size of the firm can in fact affect both the outcome and treatment participation. On the one hand, larger firms may have specific departments to perform basic research (and not only applied research) and achieve more original output. On the other, larger firms could have a higher probability of winning procurement contracts and hence to file patents induced by public demand. ix) *Backward citation lag* measures the average lag of the citations made by a patent to the prior art. As suggested by Fabrizio (2009), patents building on more recent prior art might come from technology fields or subfields experiencing more

Table 4: Probit regressions results

Originality	Coeff.	SE	Z
App-grant _lag	0.034***	0.004	8.95
# codes	0.003**	0.001	2.46
log_# claims	0.050***	0.006	8.64
log_# cites_rec	0.002	0.005	0.40
log_# cites_made	-0.060***	0.006	-9.83
log_pat_stock	0.011***	0.001	7.74
back_cites_lag	0.010***	0.001	7.08
App_Year_Dum	Yes		
HJT_subcat_Dum	Yes		
_cons	-2.625***	0.035	-74.64
N	946,721		

Novel	Coeff.	SE	Z
App-grant _lag	0.031***	0.004	8.52
# codes	0.003**	0.001	2.35
log_# claims	0.046***	0.005	8.45
log_# cites_rec	-0.003	0.005	-0.61
log_# cites_made	-0.048***	0.005	-9.68
log_pat_stock	0.011***	0.001	7.93
App_Year_Dum	Yes		
HJT_subcat_Dum	Yes		
_cons	-2.550***	0.030	-85.78
N	1,039,052		

Breadth	Coeff.	SE	Z
App-grant _lag	0.027***	0.004	6.658
# codes	0.002	0.001	1.387
log_# claims# claims	0.048***	0.006	7.778
log_# cites_rec	-0.005	0.005	-0.974
log_# cites_made	-0.043***	0.006	-7.634
log_pat_stock	0.011***	0.002	6.772
App_Year_Dum	Yes		
HJT_subcat_Dum	Yes		
_cons	-2.522***	0.034	-74.399
N	788,879		

rapid advance, and could then be more original and novel than patents building on old technologies[23].

As described in section 4.1.4, since three out of the five dependent variables are defined for different patent subsamples, I run 3 different probit regressions to properly estimate the propensity score: the first one for the *Originality* dimension, the second for the *Novel* variable, and the third for the 3 *Breadth* -related outcomes. Because the variables *Number of claims*, *Number of citations made*, *Number of citations received*, and *Patent stock* are distributed in a heavily skewed way, they are considered in logs in the estimations. Table 4 presents the results of the probit regressions. Clearly, since the variable *Backward citation lag* bases upon backward citations, it is not defined for patents that make no references to prior art. Because such patents belongs to the subsamples I use to compute the novelty-related outcomes, the variable *Backward citation lag* is used as a covariate only for estimating the propensity score that will be used to retrieve the ATT for the originality dimension.

Once that the propensity score has been recovered I proceed to the exact matching on a subset of

[23]Note that variables reporting the number of citation also takes the value zero. To take them in logs I hence add one citation to every patent, as is commonly done in the literature

covariates that I consider of major importance in determining participation in the treatment. As briefly mentioned above, since patents are considered as treated if they are induced by public procurement contract, the trend and the composition of public demand fundamentally affects the likelihood of a specific patent to be treated. For instance, a patent applied for in 2002, belonging to technology class 'Food or Edible Materials (class 426)', has completely different odds of being induced by a procurement contract compared to a patent applied for in 1985, belonging to class 'Aeronautics and Astronautics (class 244)'. In order to be sure to avoid pairing between patents that are so distant in terms of technology space and timing, I therefore perform the exact matching on the 20 *Application Year Dummies*, and on 36 *Technological subcategory dummies*. I hence construct a 36x20 matrix, in which each patent is inserted according to its application year and to the technological subcategory it belongs on the basis of its primary classification. Inside each cell treated patents are then matched to suitable control patents exploiting the propensity score. As matching algorithm I use the nearest-neighbor approach, that pairs each treated unit with the closest non-treated unit in terms of propensity score[24]. It is worth noting that the large size of the dataset allows to reach a high level of accuracy through the exact matching on technology fields and application years without forcing bad matches on the basis of the propensity score. In each of the 36 patent classes taken into account, treated units never represent more that 3% of the whole population of patents, always ensuring enough non-treated observations from which to pick a close observation in terms of propensity score. However, to further reduce the odds of producing bad matches, I also set a caliper threshold imposing a tolerance level on the maximum propensity score distance between two paired units[25].

4.2.4. Matching quality

One of the main advantages of using matching methods relies in the fact that they have straightforward diagnostics by which their performance can be assessed (Stuart, 2010). The literature proposes several strategies to evaluate whether the matching procedure manages to balance the empirical distributions of the relevant covariates in the matched treated and control groups. Rosenbaum and Rubin (1985) suggest the implementation of a two-sample t-test for significant differences in covariate means between groups, before and after the matching. Clearly, if the quality of the matching is good, the covariates should be balanced in both groups and no significant differences should be found after the pairing (Caliendo and Kopeinig, 2008). A second methodology, also proposed by Rosenbaum and Rubin (1985), entails to compute the mean standardized bias, defined as the difference in means of each covariate between the treated and the control group divided by the standard deviation in the full treated group. The size of the bias should be compared for the treated and the control group for each covariate before and after the matching. The matching procedure should clearly reduce the size of the standardized bias and, while there is no clear threshold under which it is possible to tell the success of the matching with certainty, a bias reduction below 3 or 5 per cent is generally considered as sufficient (Caliendo and Kopeinig, 2008) to confirm the good quality of the pairing.

Table 5 reports the t-test for differences in covariate means and the standardize bias for all the covariates I used to implement the matching procedure. As the table shows, while before the matching the differences in the covariate means between the treated and the control group are almost everywhere large and significant, these differences reduce abundantly and, apart form one case, are no longer significant after the matching. Only the *Backward_average_Lag* variable, for the sample used to compute the variable *Originality*, shows a significant difference between the treated and the control group after the matching. Nevertheless, also in this case as for all the other covariates taken into account, after the matching the standardized bias is below the 3% threshold, confirming that the matching procedure

[24] For a review of the available matching algorithms and of their strength and weakness see Caliendo and Kopeinig (2008).

[25] I here follow Rosenbaum and Rubin (1985) who suggest to set the caliper option to a value that corresponds approximately to .25 times the standard deviation of the propensity scores recovered with the probit regression.

Table 5: Two sample t-test for mean differences between the treated and the control group before and after matching

ORIGINALITY	NOT MATCHED				MATCHED			
	Non-treated	**Treated**	**Diff.**	**% bias**	**Non-treated**	**Treated**	**Diff.**	**% bias**
App_grant_lag	2.225	2.250	-0.025**	2.4	2.210	2.239	-0.029	2.4
# codes	4.507	4.567	-0.060*	1.8	4.562	4.567	-0.005	0.2
log_# claims	2.644	2.699	-0.054***	7.0	2.692	2.699	-0.007	0.91
log_# cites_rec	1.573	1.535	0.038***	3.3	1.554	1.536	0.017	1.5
log_# cites_made	2.270	2.178	0.092***	12.3	2.179	2.179	0.000	0.0
log_pat_stock	5.913	6.304	-0.391***	13.9	6.334	6.302	0.032	1.1
Back_avg_lag	7.420	7.568	-0.148***	4.4	7.668	7.569	0.099*	2.9
App_year dum	-	-	Yes	Yes	-	-	0	0
Tech_Class_dum	-	-	Yes	Yes	-	-	0	0
N	938,073	8,648			8,636	8,636		

NOVELTY	NOT MATCHED				MATCHED			
	Non-treated	**Treated**	**Diff.**	**% bias**	**Non-treated**	**Treated**	**Diff.**	**% bias**
App_grant_lag	2.195	2.223	-0.028*	2.4	2.197	2.212	-0.015	1.2
# codes	4.493	4.523	-0.030	0.9	4.502	4.524	-0.022	0.7
log_# claims	2.616	2.669	-0.053***	6.8	2.667	2.669	-0.003	0.3
log_# cites_rec	1.561	1.532	0.029*	2.5	1.530	1.533	-0.003	0.3
log_# cites_made	2.109	2.020	0.089***	10.1	6.358	6.301	0.057	0.1
log_pat_stock	5.900	6.314	-0.413***	14.7	6.363	6.311	0.052	1.9
App_year_dum	-	-	Yes	Yes	-	-	0	0
Tech_Class_dum	-	-	Yes	Yes	-	-	0	0
N	1,029,528	9,524			9,511	9,511		

BREADTH	NOT MATCHED				MATCHED			
	Non- treated	**Treated**	**Diff.**	**% bias**	**Non-treated**	**Treated**	**Diff.**	**% bias**
App_grant_lag	2.219	2.233	-0.014	1.2	2.217	2.222	-0.005	0.4
# codes	5.194	5.178	0.016	0.5	5.195	5.179	0.016	0.5
log_# claims	2.618	2.672	-0.054***	6.9	2.673	2.672	0.000	0.0
log_# cites_rec	1.617	1.580	0.036**	3.1	1.586	1.582	0.004	0.4
log_# cites_made	2.107	2.035	0.071***	8.2	2.048	2.036	0.012	1.3
log_pat_stock	5.936	6.302	-0.366***	13.0	6.304	6.279	0.025	2.0
App_year_dum	-	-	Yes	Yes	-	-	0	0
Tech_Class_dum	-	-	Yes	Yes	-	-	0	0
N	781,412	7,467			7,458	7,458		

Table 6: Results of the ATT estimations

Outcome	Sample	Treated	Controls	Difference	S.E.	T-stat
Originality	Unmatched	.5844	.5387	.0457***	.0034	13.32
	Matched	.5845	.546	**.0384***	.0048	7.89
Novel	Unmatched	.7840	.7590	.0250***	.004	6.58
	Matched	.7839	.7737	**.0101***	.006	1.69
Low_Breadth	Unmatched	.6354	.6368	-.0014	.005	-0.25
	ATT	.6352	.6339	**.0013**	.007	0.17
Med_Breadth	Unmatched	.5251	.4987	.0263***	.005	4.53
	Matched	.5248	.5045	**.0202***	.008	2.47
High_Breadth	Unmatched	.4025	.3497	.0528***	.0055	9.53
	Matched	.4023	.3507	**.051***	.0079	6.50

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

is doing a good job in reducing the selection bias. It is also worth noting that, since I performed an exact matching on the year of application and on the technology class, the difference in the means for the application year and technology class binary variables is evidently equal to zero everywhere after the matching implementation.

5. Results

Since the diagnostic strategies confirm the good quality of the matching procedure, I can then recover the ATT through the estimator described in section 4.2.2. Table 6 displays the results and, in particular, reports the average outcome for the treated and the control group and the difference between the two groups both before (unmatched sample) and after the matching (matched sample), for each of the five outcome variables. While the gap in in the unmatched sample clearly includes the selection bias and might hence be inflated or deflated, the differences in outcomes between the treated and the control group in the matched sample identifies the ATT and is then the parameter of interest (in bold in the table).

As the table shows, patents induced by a public procurement contract between the US Federal Government and an American private corporation exhibit a significantly higher degree of originality, as measured through the *Originality index*, with respect to control patents. In particular, procurement-induced patents have on average an originality index that is 3.8 percentage points higher than patents in the control group. This result confirms *Hypothesis 1*, presented in section 3: patents induced by public procurement contracts are indeed more original and appear to base on broader technological roots with respect to the patents the private sector would have produced in the absence of public demand. To evaluate whether or not the average treatment effect upon the *Originality index* I retrieved is entirely driven by few outliers and is actually affecting the whole distribution of the index, Figure 4 displays the kernel density distribution of the *Originality Index* for the matched sample by treatment status.

As the figure shows, the treatment substantially shifts rightward the density distribution for the patents induced by procurement contracts compared to controls, confirming the relevance of the treatment effect in affecting the whole distribution of the index for the treated patents.

Figure 4: Kernel density distribution of the Originality index by treatment status for the matched sample

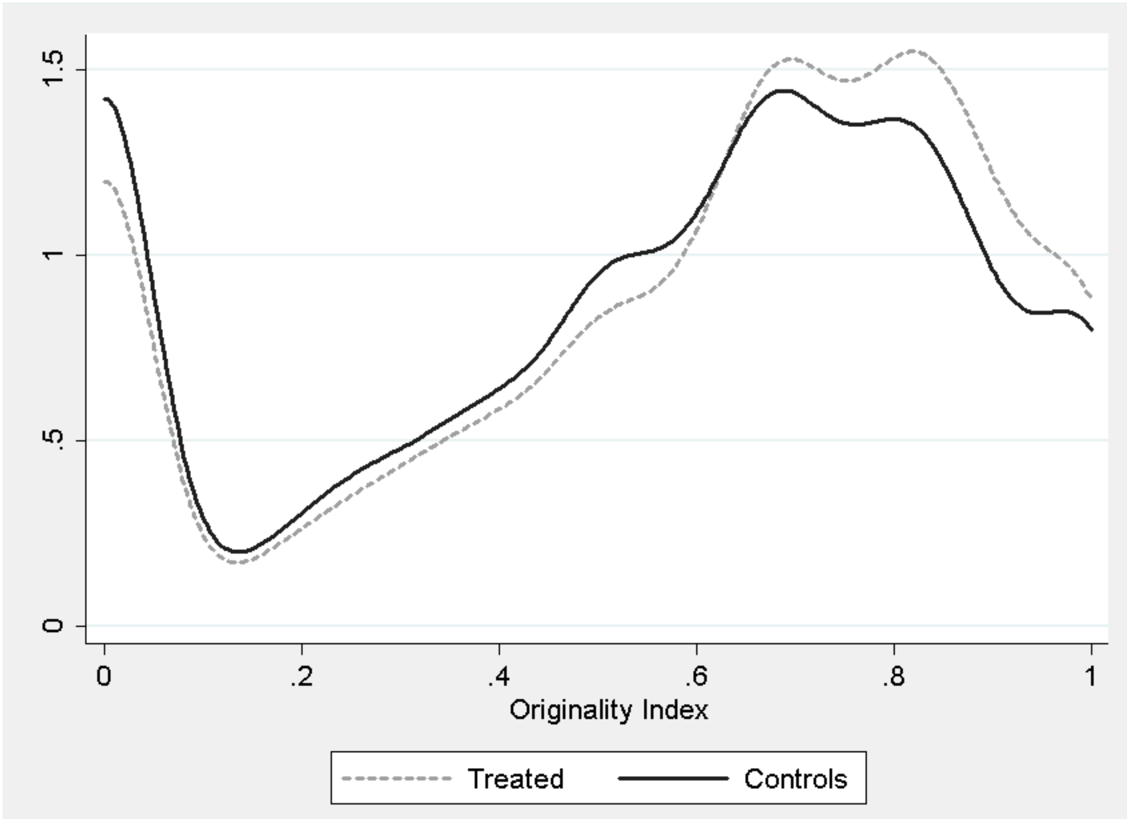

Before analyzing the results for the novelty-related measures, it should be noted that, given the binary nature of these variables, the average outcomes should be interpreted as participation rates, reporting the share of patents embodying a novel or a broad combination of technology codes. The ATT then represents the difference in the share of patents embodying a specific combination, between the treated and the control group. In the case of the *Novel* outcome variable, there is hence a positive and significant difference in the share of patents embodying novel combinations of technology codes between the treated and the control patents. More specifically, there are 1 percentage points more patent with novel combinations of codes among the patent induced by public procurement contracts with respect to control patents. Even if the difference appears to be quite small in terms of magnitude, it should be noted that, given that most patents (76% in the full sample) embody novel combinations, the treatment decreases of 4,4 % the share of the patents arising from refinements of old combinations of technology codes. This result therefore corroborates *Hypothesis 2a*, stated in section 3. Procurement related patents seems to stem from more exploratory research and to incorporate a higher share of novel combinations (i.e. less refinements) compared to the proportion that would be observed in the absence of federal procurement.

The results also confirms to some extent *Hypothesis 2b*. While there is no significant difference between the treated and the control group in the share of patents embodying broad novel combinations, as measured at the patent class level (3-digit, 470 classes, i.e. *Low_Breadth* equal 1), the situation sharply changes when we consider larger distances in terms of technology space. Both for the variables *Med_Breadth* and *High_Breadth* there are positive and significant differences in the share of patents with broad combination as measured at the 2-digit and the1-digit level. In particular there are about 2 percentage points more patent embodying novel combinations that mixes technology codes from different HJT sub-categories (2-digit, 37 cells,) and 5,1 percentage points more patents embodying novel combinations that mixes technology codes from different HJT categories (1-digit, 6 cells). Therefore a group of patent embodying novel combinations induced by public procurement contracts will include on average 4 % more patents incorporating broad combinations as measured at the 2-digit level and 14,5 % more patents including broad combinations as measured at the 1-digit level, compared to the counterfactual situation in which those patents were the output of a fully privately prompted research.

The results presented in this section hence confirm that federal procurement is able to stimulate innovations that are peculiar objects in the technology space and not just a reproduction of the same kind of innovative output that corporate research would bring about in the absence of public demand. First, the result on the originality dimension supports the idea that procurement-related patents build on more spread and diffused technological roots than corporate patents obtained through fully private efforts. This implies that procurement contracts stimulate R&D strategies that rely less on the exploitation of limited regions of the technology space and more on wider exploration of the technological landscape. In the same way, the result for the *Novelty* dimension confirms that the output of the R&D related to federal procurement contracts is less based on refinements of already exploited combinations of technological capabilities and more on new linkages between previously unconnected region of the technology space. Especially noteworthy are the findings about the breadth-related dimensions. These results not only corroborate the idea that innovative public procurement contracts spur innovations embodying more broad recombinations of existing technologies compared to the ones that the private sector would achieve by itself, but also that innovative public procurement becomes more and more crucial in spawning these broad recombinations as the technological distance between the component used in the recombination increases. The latter result fits well the idea proposed in section3 that procurement might be able to absorb part of the uncertainty connected to broader exploration strategies. Since the larger the distance between the technological capabilities, the more uncertain in term of economic and technical feasibility is the output resulting from their combination, innovative procurement might be then especially important to reduce the high uncertainties and ambiguities that private firms face in generating innovations connecting very distant regions in the technological landscape.

5.1. Robustness check: time stratification

As it has been pointed out throughout the paper, the time dimension is pivotal in many ways in the analysis and may affect both the identification strategy and the outcomes. I hence implement a robustness check that entails the stratification of the data in 4 time-windows of 5 years each: 1985-1989,1990-1994, 1995-1999, and 2000-2004. I then adopt the same empirical strategy used in the general case and described in section 4, performing 'hybrid matching' for each of the outcome variables and each time-window (i.e. 12 matching procedure[26]). The rational for running this robustness check is mainly to rule out the concerns that specific time periods might be driving the results presented in the previous section. There are several reasons why that could be the case and why it might represent a serious issue. In the first place, the period I am taking into account also includes years from 1985 to 1991, a period in which US military expenditure were still very high due to the final phase of the cold war with USSR. Figure 5 [27] clearly shows how military expenditure sharply declined during the 1990's both in absolute terms and as share of the total government expenditure.

Patents induced by public procurements during the 1980's might hence be extremely different from the others in the sample and might be the ones driving the results both on the originality and the novelty-related measures. Secondly, as I briefly discussed in section 4.1, even if I removed patents applied in the last 2 years for which I have data available (2005 and 2006), truncation effects may still affect the results since more original and novel patents tend to require more time to be granted.

[26]The analysis about the matching quality for the robustness check is available upon request
[27]Elaboration based on the SIPRI database, available at

Figure 5: U.S. MIlitary expenditures in absolute 2011 dollars and as a share of total government expenditure

In the very last years taken into account I might then lose the most original patents or, even worse, I might lose more original patents in a disproportionate fashion between the treated and the non-treated group, biasing the results. Finally, as described in section 4.1.2, the compliance with the FAR on intellectual property rights in federal procurement contracts increased over time (Rai and Sampat, 2012). For this reason, patents reporting the government interest statement may represent a smaller share of the total number of patent that should include it in the first years taken into account than in later periods. Especially problematic would be the case in which the compliance was highly correlated with the originality and the novelty of the patent. Also in this case it is therefore very important to rule out the possibility that the early period is entirely driving the results. Table 7 reports the results of the estimation achieved with the time-window stratification.

As the table shows, both the results on the originality dimension and on the *High-Breadth* variable appear to be robust and significant over time. In the case of the originality index the magnitude of the effects is larger in the first and in the last time-window, with a decline in the central ones. Also for the *High-Breadth* outcome the results show a higher impact in terms of magnitude in the first period, that lasts in the second one as well. However, I still find a decline during the 1990's, though this time limited to the second half of the decade. As for the originality measure the early 2000's displays again an increase in magnitude. While is evident that the results for these outcome variables are not driven by a single period, the tendency in the size of the effects follows to some extent the U-shaped trend in US military expenditure. This correspondence could imply that the larger the size of public demand, the less risk-averse is the research marginally induced by public procurement contract, and hence the more novel and original the marginal patent. Even though this is not a conclusion of this paper, it could be thought as a stylized fact entailing that the diversification effect hypothesized in

Table 7: Results by 5 years time-windows

Outcome	Sample	1985-1989	1990-1994	1995-1999	2000-2004
Originality	Unmatched	.0402***	.0300***	.0486***	.0582***
	Matched	**.0661*****	**.0192***	**.0301*****	**.056*****
Novel	Unmatched	.0108	.0260***	.0340***	.0257***
	Matched	**-.0155**	**.0225***	**.0279*****	**.0116**
Low_Breadth	Unmatched	.0039	- .0027	-.0132	.0092
	Matched	**.0236**	**.0141**	**.0070**	**.0100**
Med_Breadth	Unmatched	.0390***	.0308***	.0128	.0341***
	Matched	**.0480*****	**.0379****	**.0140**	**.0115**
High_Breadth	Unmatched	.0410***	.0502***	.0438***	.0749***
	Matched	**.0617*****	**.0597*****	**.0224***	**.0450*****

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

this paper is increasing with the size of public demand. The results for the variable *Novel* displays instead an opposite trend. There is no significant difference between the treated and the control group in the share of patents embodying new combinations of technology codes in the first and in the last time window. The result found in the focal estimation is then driven by the patents applied for in the 1990's. Finally, the results for the variable *Med-Breadth* are large and significant only for the first two time windows, whereas for the *Low-Breadth* outcome, as in the previous case, the difference between treated and control patent is never significantly different from 0, no matter the time-window considered.

5.2. Robustness check: treatment by department

In section 4.1.2 I described as the government interest statement also allows to disentangle which department awarded the contract that led to the patent. As figure 1 shows, most of the patents originate from a contract with the Department of Defense (51,7%), the Department of Energy(37,2%) , and, to a lesser extent, the NASA(7,8%). In the second robustness check I exploit this variability in the treatment variable. Instead of considering as treated the patents that include the government interest statement, I create 3 treatment variables, *Treatment_Defense*, *Treatment_Energy*, and *Treatment_NASA*, based on which of the main departments awarded the contract. I then implement the same empirical strategy adopted so far (i.e. hybrid matching) for each outcome and treatment variable. The rational for performing this robustness check is very much related to the previous one. I would like to understand whether the results of the focal estimation are entirely driven by the demand coming from a single department and, in particular, if defense-related demand is playing the prominent role in determining the previous results. If that was the case, it would be harder to make general statement about the kind of research induced by procurement contracts. However, a caveat should be made because it is not possible to unequivocally identify the patents induced by contracts with the DoD as related to military demand and the patents originated by contracts with DoE as more civil-oriented. While the DoE's responsibilities go from energy production and energy conservation to radioactive waste management and environmental research, it is also in charge of the U.S. nuclear weapons program and of the nuclear reactor production for the Navy. Therefore, even if the

Table 8: Results by stratification of the treatment variable

Outcome	Sample	Defense	NASA	Energy
Originality	Unmatched	.0340***	.0193	.0573***
	Matched	**.0386***	**.0341***	**.0415***
Novel	Unmatched	.0182***	.0169	.0310***
	Matched	**.0022**	**-.00195**	**.0313***
Low_Breadth	Unmatched	.0055	-0429**	-.0053
	Matched	**.0151**	**-.0447**	**.0007**
Med_Breadth	Unmatched	.0364***	.0208	.0138
	Matched	**.0397***	**-.0023**	**.0089**
High_Breadth	Unmatched	.0427***	.0533***	.0654***
	Matched	**.0418***	**.0214**	**.0557***

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

composition of DoE demand is clearly more diversified, it is far from being entirely civilian in nature.

Table 8 displays the results of this robustness check. As the table reports, the results for the emphOriginality variable appear to be robust for all the three department-specific treatment variables.

Again also in the case of *High-Breadth* variable the difference between the treated and the control group in the share of patents embodying distant novel combinations is large and significant both for the *Defense* variable and for the *Energy* variable. The NASA variable is not significant, probably also because of the small size of the treated group in that case. Especially interesting is the fact that the result for the *Novel* variable is positive and significant only for the DoE-related treatment. Both the results in the focal estimation and in the previous robustness check are hence driven by patents induced by contracts with the DoE. It seems then that during the 1990's, exactly when the U.S. military expenditures decreased, the DoE contracts led to patents embodying a particularly large share of novel combinations of technology codes.

6. Conclusions

In recent times innovative public procurement has been increasingly acknowledged as an effective industrial policy to spur sinnovation and raise private investment in R&D both by researchers and by policy makers (EU, 2010; OECD, 2014; Warwick and Nolan, 1994). Along with political attention, a growing body of works tried to empirically assess the impact of such a demand-side policy on firms innovative behavior. Very few theoretical and empirical contributions, though, devoted attention to the nature of the technological change that innovative public procurement brings about. The present paper addressed this issue. In particular, I hypothesized that procurement-related R&D, being less constrained by binding budgets and short term profitability goals, is more exploratory in nature compared to private R&D and will result in more original and basic innovations compared to the the ones the private sector would realize in the absence of public demand.

To test this hypothesis, I made use of patent data as a proxy for the outcome of the innovation process. In particular, I created an original dataset composed by 1,039,052 patent issued by the

USPTO between 1985 and 2004, exploiting information coming from three different sources: the *NBER patent data project*, the *USPTO Patent Full-text and Image Database*, and the *U.S. Patent Grant Master Classification File*. The information included in the dataset allowed me to construct different measures to grasp the peculiarities of the innovations generated by public procurement. First, I use patent citations to compute the *Originality Index*, developed by Trajtenberg et al. (1997) and abundantly used in the literature (Thursby et al., 2009; Guerzoni et al., 2014), to assess the basicness of a patent. Secondly, to overcome some drawbacks of this measure, I built four additional novelty-related outcome variables based on the different combinations of technology codes assigned to patents and on their technological distance, as proposed by Strumsky et al. (2012) and Youn et al. (2014). I then designed a quasi-experiment exploiting the Federal Acquisition Regulation that rules intellectual property in federal procurement contracts in the United States. A patent is considered as treated if it was induced by a contract between a federal agency of the U.S. government and a private firm based in the United States. The originality and novelty-related measures of the treated group are then compared with the outcomes for control patents. To reduce the selection bias that typically affects non-experimental studies I implemented the 'hybrid matching' technique, mixing exact and propensity score matching.

The results and the robustness checks confirm the main hypotheses made in this paper. Patents induced by public procurement contracts are on average more original and embodies more novel and more broad combinations of technological capabilities compared to suitable control patents. Particularly robust are the result on the *Originality* measure and on the *High_Breadth* variable, that reports whether a patent mixes technology codes that are very distant in terms of technology space. The relevance of public procurement in spurring distant recombination of existent technological capabilities seems hence to be increasing with technological distance.

As highlighted throughout the paper, these results closely relate to and try to provide a preliminary explanation for the findings provided in Raiteri (2014). In that case, receiving a citation from a patent induced by a public procurement contract raised the degree of generality, or pervasiveness, of the cited patent. The result of this paper sheds some light on the nature of this specific citation. In an environment characterized by high uncertainty, as for example technologies in their early developments, innovations that apply the infant technology generate new knowledge that reduces the uncertainty associated with further innovation (Bewley, 2001). Procurement induced innovation therefore generate a substantially different kind of knowledge with respect to the private R&D's output. Federal procurement seems in fact to produce innovations that are peculiar objects in the technology space, building on broader technological roots and embodying more novel and wider combination of technological capabilities. They will hence reduce the costs of explorations for further innovations by providing additional knowledge about distant region in the technology space. If, as put forward by Fleming and Sorenson (2004), science may be conceived as a map of the possible areas of the technology space to be explored, procurement induced innovations may be seen as the infrastructure that actually connect two remote points in the technological landscape, reducing their relative distance and hence raising the likelihood of further innovations combining these and proximal capabilities.

Clearly this work also presents some limitations. In the first place, as every empirical study relying on patent data as a proxy for the output of innovative activities, I consider only a portion of the whole innovation universe: patentable inventions. I therefore miss all the innovations that do not meet patentability requirements and those innovations that have been kept secret. The latter issue could particularly problematic in this context since, is some case, the government may explicitly ask a contractor to keep secret an innovation induced by a procurement contract for reasons of national security. Nevertheless, even if this problem could constrain the external validity of my analysis, the patents that I identify as treated are unequivocally induced by procurement contracts.

Secondly, it should be made clear that, due to the nature of the data at hand, the methodology used in the empirical analysis only takes care of the selection bias due to observable heterogeneity among the patents.

From the policy point of view, the results of the present paper seem to suggest that a highly innovative public demand might be pivotal for achieving greater variety and novelty in the technology ecosystem. Public procurement might hence be an especially effective technology policy tool not only for raising private investment in R&D and spawning more of the same technological change that the private sector would bring about, but also for triggering exploratory searches in the technology space and fostering more basic and original innovation in an era in which innovating through novel and distant recombinations is becoming increasingly difficult.

References

Akcigit, U., Kerr, W. R., and Nicholas, T. (2013). The mechanics of endogenous innovation and growth: Evidence from historical us patents.

Antonelli, C. (1995). *The economics of localized technological change and industrial dynamics*, volume 3. Springer.

Arthur, W. B. (2009). *The nature of technology: What it is and how it evolves*. Simon and Schuster.

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*. Springer.

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.

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

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. (2009). Why challenge the ivory tower? new evidence on the basicness of academic patents*. *Kyklos*, 62(4):488–499.

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. A. (1988). *Path-dependence: putting the past into the future of economics*. Institute for mathematical studies in the social sciences.

Draca, M. (2012). Reagan's innovation dividend?

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. (2009). Absorptive capacity and the search for innovation. *Research Policy*, 38(2):255–267.

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*.

Fleming, L. (2001). Recombinant uncertainty in technological search. *Management science*, 47(1):117–132.

Fleming, L. and Sorenson, O. (2004). Science as a map in technological search. *Strategic Management Journal*, 25(8-9):909–928.

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.

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

Georghiou, L., Edler, J., Uyarra, E., and Yeow, J. (2013). Policy instruments for public procurement of innovation: Choice, design and assessment. *Technological Forecasting and Social Change*.

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

Griliches, Z., Pakes, A., and Hall, B. H. (1988). The value of patents as indicators of inventive activity. NBER Working Papers 2083, National Bureau of Economic Research, Inc.

Guerzoni, M. (2010). The impact of market size and users' sophistication on innovation: The patterns of demand. *Economics of Innovation and New Technology*, 19(1):113–126.

Guerzoni, M., Aldridge, T. T., Audretsch, D. B., and Desai, S. (2014). A new industry creation and originality: Insight from the funding sources of university patents. *Research Policy*.

Guerzoni, M. and Raiteri, E. (2013). Demand vs. supply technology policies: hidden treatment and new empirical evidence on the 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.

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.

Imbens, G. W. and Wooldridge, J. M. (2009). Recent developments in the econometrics of program evaluation. *Journal of Economic Literature*, 47(1):5–86.

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.

Kuznets, S. (1962). Inventive activity: Problems of definition and measurement. *... and direction of inventive activity: Economic and social ...*, I:19–52.

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.

Lechner, M. (2000). An evaluation of public-sector-sponsored continuous vocational training programs in east germany. *Journal of Human Resources*, 35(2):347–75.

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.

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*, 17:371–399.

March, J. G. (1991). Exploration and exploitation in organizational learning. *Organization science*, 2(1):71–87.

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

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. 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.

OECD (2013). Government at a glance 2013.

OECD (2014). Intelligent demand: Policy rationale, design and potential benefits.

Pakes, A. and Griliches, Z. (1980). Patents and r and d at the firm level: A first look.

Puhani, P. A. (2002). Advantage through training in poland? a microeconometric evaluation of the employment effects of training and job subsidy programmes. *Labour*, 16(3):569–608.

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

Raiteri, E. (2014). A time to nourish? evaluating the impact of innovative public procurement on technological generality through patent data.

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.*

Rosell, C. and Agrawal, A. (2009). Have university knowledge flows narrowed?: Evidence from patent data. *Research Policy*, 38(1):1–13.

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.

Rubin, D. B. (1973). Matching to remove bias in observational studies. *Biometrics*, pages 159–183.

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.

Slavtchev, V. and Wiederhold, S. (2011). Government demand composition and innovative behavior in industries. *Ifo Working Paper*, 35.

Strumsky, D., Lobo, J., and van der Leeuw, S. (2011). Measuring the relative importance of reusing, recombining and creating technologies in the process of invention.

Strumsky, D., Lobo, J., and van der Leeuw, S. (2012). Using patent technology codes to study technological change. *Economics of Innovation and New technology*, 21(3):267–286.

Stuart, E. A. (2010). Matching methods for causal inference: A review and a look forward. *Statistical science: a review journal of the Institute of Mathematical Statistics*, 25(1):1.

Thursby, J., Fuller, A. W., and Thursby, M. (2009). Us faculty patenting: Inside and outside the university. *Research Policy*, 38(1):14–25.

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.

Warwick, K. and Nolan, A. (1994). Evaluation of industrial policy.

Weitzman, M. L. (1998). Recombinant growth. *Quarterly Journal of Economics*, pages 331–360.

Youn, H., Bettencourt, L., Strumsky, D., and Lobo, J. (2014). Invention as a combinatorial process: evidence from us patents. *wp forthcoming.*