Do Patent Assertion Entities Harm Innovation? Evidence from Patent Transfers[1]

Gianluca Orsatti[‡,◊] and Valerio Sterzi[‡,2]

[‡] GREThA - UMR CNRS 5113, University of Bordeaux

[◊] University of Turin

November 13, 2018

[1]All our thanks to Laurent Berge, Stefania Fusco, Fabian Gaessler, Philippe Gorry, Stuart Graham, Francesco Lissoni, Catalina Martinez, Pierpaolo Parrotta, Julien Penin, David Schwartz, Tim Simcoe, Gianluca Tarasconi, Marco Vivarelli, Georg von Graevenitz and to participants in the: 2018 Workshop of Economics of Innovation and Economics of Science in Strasburg (FR), 2018 Max Planck Seminar at Munich (DE), 2017 VIII IBEO Workshop in Corte (FR), 2017 Cournot seminar at Beta in Strasbourg (FR), EPIP2017 conference in Bordeaux (FR), 2017 Via Inno Seminar at University of Bordeaux (FR), 2018 RITM seminar at University of Paris Sud (FR), 2018 CREA Seminar at University of Luxembourg (LU), 2018 EPIP-GREThA Workshop in Bordeaux (FR), 2018 Druid Conference in Copenhagen (DK). We also acknowledge the Agence Nationale de la Recherche (ANR) for its financial support (NPEIE Project, no. ANR-17-CE26-0014-01).

[2]Corresponding author: Valerio Sterzi, Avenue Leon Duguit, F-33608 Pessac Cedex. Tel: +33-556844051. Email: valerio.sterzi@u-bordeaux.fr

Abstract

The recent upsurge of patent litigation cases initiated by patent assertion entities (PAEs) in the US has led to an intense debate about their effect on innovation dynamics and on the IP system functioning. We contribute to this debate by providing original evidence on high-tech patent transfers involving PAEs at the European Patent Office from 1997 to 2012. Our evidence is threefold. First, relative to both patents transferred to producing companies (PEs) and never-transferred patents, patents purchased by PAEs are on average of higher technological quality. PAEs may thus increase liquidity in the patent market, enhancing its efficiency. Second, around the transfer and in the post transfer period, the citation profile of patents transferred to PAEs falls considerably, suggesting that PAEs do not play as intermediaries in the market. Interestingly, also patents transferred to PEs do receive fewer citations after the transfer, suggesting an increasing incidence of strategic patent acquisitions in the ICT domain and opening relevant questions about its entire functioning. Finally, a finer comparison of citation profiles between patents transferred to PAEs and patents transferred to PEs shows a non significant difference in the post-transfer citation drop: also PEs, when involved in transfers of specific technologies, seem not to purchase patents for inventing around the protected technologies.

Keywords: Non-practicing entities; Patent assertion entities; Patent trolls; Patent transfers; Patent citations, Market for technology.

JEL Codes: O31, O34, D23

1 Introduction

Once seen merely as a means of protecting an invention, patents are now considered as marketable assets that can be acquired, held, licensed and sold strategically (Papst, 2012). Markets for technology have expanded rapidly in the last 20 years or so. According to Ocean Tomo (Elsten and Hill, 2017), in 2015 intangible assets (mainly patents, software, trademarks and copyrights) represented 84% of the S&P 500 market capitalization – corresponding to 16% of growth from 1995 – and 71% of that of the S&P Europe 350.

Due to increased opportunities for patent monetization, the activity of companies that facilitate the transfer of exclusive rights to inventions has recently experienced a tremendous upsurge (Hagiu and Yoffie, 2013). Consequently, new intermediaries such as patent aggregators and patent assertion entities (PAEs) have become quite influential and controversial, especially in the ICT industry.

A politically diffused opinion is that patent trolling[1] is becoming a growing concern (Cohen et al., 2016; Lemley and Feldman, 2016) or even the "most significant problem facing the patent system today" (Lemley, 2006, p. 2). Over the past decade, the US patent system has indeed experienced an explosion of litigation cases initiated by PAEs. Recent studies estimate that the PAE business in the US is worth around $30 billion in settlements and licensing fees annually (Carter, 2013; Yeh, 2013).[2] Not surprisingly, a heated debate has intensified on the economic role that these companies play in the market for patents and on their impact on innovation dynamics. Indeed, in reaction to the proliferation of patent lawsuits initiated by PAEs, the US Congress recently introduced several bills proposing to finely regulate the process of patent licensing and assertion. The new inter partes reviews implemented by the 2011 American Invent Act and a number of subsequent U.S. Supreme Court decisions over issues such as patentable subject matter, attorney fees and forum shopping have been directed to curtail the PAEs' activity (Fusco, 2016).

[1]PAEs are sometimes called, in derogatory terms, "patent trolls".

[2]In 2016, about 67% of all US patent lawsuits were filed by non-practicing entities (the large majority represented by PAEs), up from the 61% experienced in 2015 (2015 Patent Dispute Report, Unified Patents; figures available at `https://www.unifiedpatents.com/news/2016/5/30/2015-patent-dispute-report`

To study the impact of PAEs on innovation, most of the existing empirical research has focused only on patent litigation, leaving out of the analysis the enforcement activities settled out of court, *i.e.* those that did not become public.[3] However, figures based exclusively on patent litigation give a very partial intuition of the importance of PAEs' activities in the patent market, together with their implications, as *"these visible actions are just the tip of the iceberg"* (Shapiro and Scott-Morton, 2014). Indeed, instead of going through litigation, PAEs are more likely to prefer to set royalty demands strategically below litigation costs in order to make the business decision to settle an obvious one (Leslie, 2008). This behavior makes it difficult to trace their business and to properly analyze their impact, given that data only based on patent litigation provide evidence on selected targets and underestimate the PAE presence in the market for patents.

This paper contributes to the literature by building a unique database of PAEs' patenting history and patent transfers at the European Patent Office (EPO). This allows for a wider and more systematic identification and analysis of the activity of PAEs in Europe, overcoming part of the limitations related to patent litigation data and providing original evidence on a region where PAEs are increasingly active.[4] While it is true that patent monetization is relatively less often pursued in Europe compared to the US, due to a combination of fragmentation of intellectual property jurisdictions and smaller damages awards (Mayergoyz, 2009), PAEs nonetheless increasingly account for a substantial and largely unrecognized share of patent litigation in Europe (Fusco, 2013; Ortiz, 2016).[5] Indeed, recent figures demonstrate

[3]Exceptions are represented by Fischer and Henkel (2012) and Leiponena and Delcamp (2018).

[4]As an example, Technicolor sold its patent licensing business to Interdigital early this year (https://www.technicolor.com/news/closing-sale-technicolors-patent-licensing-business-interdigital). The deal involved approximately 18,000 patents and applications across a broad range of technologies.

[5]Thumm (2018) provides an in-depth discussion of the main reasons why the PAEs' activity has recently increased its focus on the European market. On the one hand, recent patent reforms, and in particular the America Invents Act, reduced the opportunities of asset monetization for PAEs in the US. At the same time, several recent US court decisions have set legal precedents that both limit the likelihood of obtaining an injunction and make it harder to acquire and assert software-related patents. Conversely, EU institutional and legal changes and the imminent introduction of the Unified Patent Court (UPC) and the Unitary Patent (UP) are making the European patent monetization

that their presence in European courts is not negligible: a recent study of Darts-IP (2018) shows that during the period 2007-2016 PAE-related litigation in Europe grew about 19% year by year.[6]

The aim of the paper is to provide a broader understanding of the PAEs' activity and its impact on innovation, going beyond litigation cases and looking directly at patent transfers.

To assess the impact of PAEs on innovation, we look at the pattern of citations received by high-tech granted patents filed by practicing entities (PEs)[7] and subsequently acquired by PAEs. The idea is that (forward) citations are an indicator of the use of the protected technology by innovating and producing companies: patents frequently cited are patents that constitute important prior art for further related technological improvements; conversely, patents that stop being cited are patents whose technological utility reduces.

In the first part of the analysis, we compare citations received by patents acquired by either PAEs or PEs, before and after the transfer, with citations received by never transferred patents. Our first finding is that patents acquired by PAEs are, on average, of high technological quality. Before a transfer takes place, their citation rate is indeed significantly higher than the citation rate of both never transferred patents and patents transferred to PEs (around 11% more citations per year). However, after a transfer occurs, PAE-owned patents show a dramatic citation drop relative to the other groups (around 10% fewer citations per year).

To reduce selection issues, we then perform propensity score matching analyses to match patents that are transferred (either to PAEs or to PEs) with those that are never transferred on observable fixed patent characteristics. Moreover, we also include patent fixed effects in our empirical models to further control for fixed unobservable patent characteristics. Results of this second part of the analysis confirm the main findings mentioned before: PAE-acquired patents are of high technological quality and show a citation

landscape potentially more attractive for PAEs.

[6]In this regard, a coalition of companies (IP2Innovate) including among others Adidas, Daimler, Intel, Google, SAP and Spotify has urged the European Commission (Reuters, Apr 5, 2017) to take action against the explosion of lawsuits brought in Europe by PAEs.

[7]The term "practicing entities" is here used as opposite to PAEs and refers to all the entities that are not PAEs.

drop in the post-transfer period. Interestingly, we find that also patents transferred to PEs show a citation drop in the post-transfer period. This last result suggests that PEs operating in high-tech industries largely purchase patents for reasons that do not necessarily conceive the internal use of the protected technologies.

As a final step of the analysis, we strictly focus on the group of transferred patents, so that we can directly compare patents acquired by PAEs with those acquired by PEs (here, again, we perform propensity score matching techniques to compare patents that are similar on observable fixed, pre-transfer characteristics). Results reveal, on average, a negative impact of the transfer on the rate of forward citations. However, we do not find any significant difference between the two groups, suggesting that both PAEs and PEs acquire patents for reasons that do not deal with the further direct technological usage of the protected inventions. PAEs do not behave as intermediaries but seem to not create an additional obstacle to innovation by further discouraging producing companies from entering and investing in fields related to the transferred patent.

The rest of the paper is organized as follows. In the next section we discuss the theoretical background. Section 3 introduces the various data sources we use in our analysis and provides some descriptive evidence about the PAE activity at EPO. Section 4 describes the empirical strategy and the main variables we decide to implement. Section 5 presents the results and Section 6 concludes. Finally, the Appendixes present a finer description of the methodology we implement to build our final database, the tests we perform about the matching strategies adopted in the empirical analysis, and various further robustness checks to our empirical results.

2 Theoretical background

Do PAEs affect innovation? The rise of PAEs has sparked a debate regarding their value and impact on innovation. The main point of contention is whether patent enforcement pursued by these entities is an efficient mechanism for technology transfer and the creation of new products, or whether it is simply a means of collecting money for avoiding litigation (constituting a hidden cost for innovators, thus reducing incentives to perform R&D). The answer matters not just for the debate over the desirability of the existence of PAEs but, pragmatically, for guaranteeing long-run rates of technological diffusion and the efficiency of the patent system as a whole.

On the one hand, advocates of PAEs argue that such entities, by acting as intermediary organizations and helping financially constrained inventors to enforce their patent rights, enhance the efficiency of the market for inventions. On the other hand, due to the fact that a threat of legal action is sufficient to receive damages or settlement payments, regardless of actual patent infringements, opponents of PAEs argue that these entities simply exploit imperfections in the market for patents, extracting unjustified rents from producing and innovating firms.

2.1 PAEs as market-makers

The patent market consists mainly of bilateral transactions, either sales or cross-licenses, between technology suppliers and potential buyers interested in developing a particular technology. Companies privately negotiate such deals, sometimes involving hundreds or thousands of patents.[8] Outside of these bilateral deals, patent buyers and sellers frequently have a hard time finding each other, since searching for and identifying potential partners requires considerable time, effort and competences. The patent market is indeed characterized by information asymmetries on both sides: given the embryonic nature of innovation processes, knowledge suppliers have better knowledge of the intrinsic value and characteristics of their inventions; while buying companies can better evaluate the commercial value of those inventions. Likewise, the technological value of an invention is subject to strong complementary and portfolio effects (Gans and Stern, 2010; Parchomovsky and Wagner, 2005). In this sense, patent intermediaries may serve to connect those who have inventions with others who can create products from the inventions (Khan, 2013) and may also strengthen demand within IP markets by offering a viable “exit” for innovator who are looking for ways to extract value from patents by means other than practicing (Papst, 2012; Serrano and Ziedonis, 2018).

[8]For example, in June 2011, a consortium of Apple, Microsoft, Sony, and several other large technology companies outbid Google to buy Nortel’s 6,000 patents and patent applications for $4.5 billion.

Moreover, the asymmetry in financial resources between small inventors and large patent holders and manufacturers prevents the former from making a credible threat to litigate against infringement (Haber and Werfel, 2015). This is due to the high costs associated with litigation[9] (especially in cases of defeat in court) and to a lack of resources, time and know-how, on the inventor's side.[10]

The combination of high search costs and financial constraints paves the way, in principle, to intermediaries between inventors and investors, thereby providing the opportunity to economize on the costs of expertise to identify and sell profitable inventions (Lizzeri, 1999; Hoppe and Ozdenoren, 2005). According to this view, if PAEs behave as intermediaries, they may improve the efficiency of the market for technologies, indirectly spurring innovation. PAEs may thus act as intermediaries that identify undervalued patents and invest time and resources to find other firms interested in those patents (McDonough, 2006).

In all, the two main arguments in favor of the PAE business are the following: (i) PAEs provide inventors with competences, capital, and bargaining power, enhancing the incentives to innovate; and (ii) PAEs serve an intermediary function in the patent market by connecting patent holders with entities that can create profits from their inventions, increasing the efficiency of the market for technologies.

2.2 PAEs as market-breakers

A conflicting view suggests instead that the main PAE business is to extract unjustified rents from productive and innovative firms. These extra rents originate from the inefficiency of the legal patent system (Burk and Lemley,

[9]According to Lanjouw and Schankerman (2004) small patentees are relatively disadvantaged in enforcing their IPRs and thus more likely to litigate than negotiate.

[10]For example, France Brevets, the sovereign patent fund established by the French government, has the mission to help small and medium French companies and public research centers to monetize their patent portfolios. In 2011 France Brevets signs an agreement with Inside Secure, a French company specialized in secure transactions, for the exclusive license of 70 NFC (near field communication) patents. Two years later, France Brevets files patent infringement lawsuits against HTC and LG in the US and in Germany for using two patents (US 6700551; US 7665664) that were granted to Inside Secure in 2004 and 2010. LG decides to settle in 2014, while HTC does not but loses the patent litigation case in 2015.

2009; Feldman, 2012), where a threat of legal action is sufficient to induce targeted firms to settle, regardless of the actual patent infringement (Lemley and Shapiro, 2006). Whenever a patent holder can obtain an injunction that will force the downstream producer to take the product off the market, the threat can be very effective.

This is particularly true for complex technologies and, in general, for all inventions in the information technology sector in which many patents are possibly associated with a single product and, particularly, when manufacturers have already invested irreversible technology-specific capital (Lemley and Shapiro, 2006). Since PAEs do not depend on the final product market, conventional market remedies, i.e. cross licenses, are ineffective in preventing PAEs from pursuing holdup strategies (Lu, 2012).

In order to extract licensing fees, PAEs often engage in frivolous litigation (Lu, 2012; Feng and Jaravel, 2016), imposing litigation and licensing costs that are not proportionate to the value of the patented technology, thereby creating an unwanted tax on innovative products and services (Feldman and Frondorf, 2015). In this case, the PAEs presence in the market augments the risk for producing companies of being sued. This unwanted and inefficient extra cost may have an indirect, negative impact on innovation, inducing firms to reduce or even interrupt their R&D investments and to shift focus in order to avoid future litigation.[11]

2.3 Main evidence

Theoretical studies reveal potential negative impacts of PAEs on innovation dynamics (Lemley and Shapiro, 2006; Reitzig et al., 2007; Turner, 2011; Penin, 2012). This is coherent with anecdotal evidence (Cohen et al., 2016). However, empirical evidence about the consequences of PAEs on innovation is rather inconclusive. Importantly, the extant literature has mainly focused on the direct impact of PAEs on targeted firms in terms of additional licensing and extra litigation costs to sustain, while the indirect consequences on the market for innovation, taken as a whole, have not been deeply investi-

[11] This explains the increasing importance of defensive patent aggregators, such as RPX and AST, which purchase patents to mitigate the risk and the cost of litigation, offering a sort of insurance against patent troll risk to inventors and producing companies (Papst, 2012; Hagiu and Yoffie, 2013).

gated.

An important shortcoming of the extant evidence is that it is mainly based on patent litigation data. These data have been used by a number of legal scholars and economists mainly to (1) find evidence of "opportunistic" behavior of PAEs and to (2) evaluate the impact of litigation on R&D investments and sales of innovating companies targeted by PAEs.

With regard to the first point, results are mixed. Some authors suggest that PAEs behave opportunistically. Feldman and Frondorf (2015) surveyed the in-house legal staff of 50 product companies characterized by initial public offerings (IPOs) between 2007 and 2012. They found that 40% of respondents received patent demands during the time of their IPOs, with those demands coming mainly from PAEs. Cohen et al. (2014) found that cash availability is the principal determinant of PAEs' litigation targeting, while this is not true for small inventors and producing companies. Love (2013) found that PAEs litigate their patents late in the patent life, waiting until a lucrative industry has developed before filing suit. Feng and Jaravel (2016) found that PAEs purchase more patents that are "more obvious and contain vaguer claims", suggesting that they acquire patents with the sole purpose of litigation.

While it is true that PAEs target successful commercializers and cash-rich firms, this does not necessarily imply that their litigation are as "frivolous" as suggested by the anecdotal evidence. Indeed, recent works found that PAEs are not (mainly) involved in frivolous litigation and, interestingly, they do not seem to assert low-quality patents. As selected examples, Shrestha (2010), comparing a sample of patents litigated by 51 PAEs to a sample of patents litigated by other entities, found that the former were of higher quality (i.e. more cited and with a wider technical breadth). Risch (2012) analyzed the patents asserted by the ten most-litigious PAEs in the US and found them to be qualitatively similar to those asserted by producing companies. Similarly, focusing on patents acquired (instead of patents litigated) by PAEs, Fischer and Henkel (2012) and Leiponena and Delcamp (2018) found evidence suggesting that PAEs acquire patents of high technological quality.

With regard to the second point, the extant literature substantially agrees that the (litigation and licensing) costs to targeted firms are high and that reductions in R&D and other investments are quite substantial (Cohen et al., 2014). For example, Tucker (2014) examined a case study on how the actions of Acacia Research Corporation, a well-known PAE,[12] have affected technology sales of US firms in the field of medical imaging technology. She found that sales of products protected by patents affected by litigation with Acacia have considerably diminished as a consequence of a reduction in incremental product innovation during the period of litigation. Bessen et al. (2011), analyzing the defendant's stock market events around the filing of patent lawsuits involving a PAE over the period 1990-2010, found that these lawsuits were associated with half a trillion dollars of lost wealth to defendants. Finally, Bessen and Meurer (2013) estimated the direct costs of defendants in litigation with PAEs at about $29 billion in 2011.[13]

If PAEs do impose high costs on the targeted firms, it is however possible that they serve as tax collectors for inventors from whom patents have been bought. Payments from innovative companies might not be considered as a reduction in R&D efforts if they are counterbalanced by significant transfers to the original inventors. However, early evidence is not encouraging. Bessen and Meurer (2013) used survey evidence on US companies and found that payments to independent inventors only account for 5% of the direct costs that defendants incur in litigation with PAEs, while 62% goes to PAEs' operating costs (including 15% which goes to payments to the PAEs' own R&D departments), 23% to legal expenses, and 10% to profits.

In this paper we take an original perspective in investigating the effect of PAEs on innovation. By looking at patent transfers, we indeed empirically test for the effect of a patent transfer to a PAE on the further usage of the acquired technology.

Our claims are straightforward: if PAEs behave as intermediaries and perform the role of creating new opportunities for technological development, we do expect to see them selecting particularly valuable technologies and, after acquiring them, to find a better positioning of those technologies into

[12] Quinn (2010) labeled Acacia as the "mother of all patent trolls".

[13] However, Schwartz and Kesan (2013) contested the analysis proposed by Bessen and Meurer (2013), arguing that their results are not based on a random or representative sample, and that the $29 billion cost estimated by Bessen and Meurer (2013) should be viewed as the "highest possible limit".

the market, enhancing their usage.

Conversely, if the PAE business is mainly related to collecting rents from producing companies through the threat of legal actions, we do expect the opposite to emerge. Indeed, if this is the case, PAEs do not target technologies for their intrinsic value but for their possible enforcement and, importantly, the absence of intermediary actions should flatten the innovation activity around the technologies they buy, reducing their further usage.

3 Data

To build a database of patent transfers involving PAEs (and PEs) at the European Patent Office (the "PAE-EP database"), we first produce an extensive list of PAEs active in the European technology market. We do so by exploiting several external sources of information about PAEs that are active worldwide. Then we match the PAE list with the list of applicants retrieved from the EP-Register database[14] to track their patenting history at the EPO.

3.1 Database construction

The PAE list We broadly define PAEs as independent organizations (legal entities) which own or purchase patents filed from or granted to other companies or individual inventors without the intent of developing, producing and/or commercializing the related products or processes. In most cases, these firms do not conduct any R&D activity. Their main business consists in generating revenues by asserting acquired patents against alleged infringers (Chien, 2008). This definition excludes certain inventors that are often considered as non-practicing entities, in particular individual inventors, universities and academic institutions who initiate suits.[15]

To individuate active PAEs, we exploit multiple sources. As a primary source of data we collect information contained in patent litigation data

[14] https://www.epo.org/searching-for-patents/legal/register.html#tab-1

[15] Wisconsin Alumni Research Foundation and Virginia Tech Intellectual Properties are two examples of academic institutions that are used to initiate patent suits. For this reason they are often labeled as PAEs. However, due to their academic nature, we decide to not consider them as PAEs.

from the UK, Germany and the US. We put together PAEs' names originally collected by Love et al. (2017) [16] [17] and by Cotropia et al. (2014).

We then complement this list of PAEs with information from web sites specialized in monitoring the PAE activity. Twenty-five PAEs active in the European market for patents are retrieved from PatentFreedom, a for-profit organization that gathers and analyzes data about PAEs.[18] A second source of data comes from IP-Checkups, a web resource that extensively collects names of active non-practicing entities worldwide. More precisely, IP-Checkups provides a partial list of eleven PAEs, together with a comprehensive list of related subsidiaries.[19]

By making use of these diverse sources, we end up with a final list of potentially active PAEs, composed of 6,127 unique entities.[20] After applying the matching procedure described below, we identify 1,752 unique entities effectively operating in the European market for patents (i.e. owning at least one EP patent). This number reduces to 1,047 when we assign subsidiaries to the main companies if the information is available.

The European Patent Register To build a unique database of European patents owned by PAEs we rely on information provided by the European Patent Register (EPR, November 2015). The EPR contains all the publicly available bibliographic, procedural and legal status information on European patent applications as they pass through each stage of the granting process. More precisely, as highlighted by the European patent system

[16]Love et al. (2017) define 7 groups of potentially non-practicing entities: (1) IP Licensing Co., Acquired Patents; (2) IP Licensing Co., Owned by Inventor or Failed Product-Producing Co.; (3) University, University IP Licensing Spin-off, or Other Research Institution; (4) Start-up, Suing Pre-Product; (5) Individual; (6) Industry Consortium; (7) IP Subsidiary of a Product-Producing Co. For the purpose of our study, we only extract information contained in groups (1) and (2).

[17]We thank Fabian Gaessler for providing these data.

[18]The names of PAEs are reported by Fusco (2013).

[19]`http://www.ipcheckups.com/npe-tracker/npe-tracker-list/`.

[20]Most of them are subsidiaries or ad hoc companies that appear to have been formed solely to hold and enforce a patent or a small portfolio of patents. We rely on IPcheckup (`https://www.ipcheckups.com/blog/a-list-of-some-npes/`) and Plainsite.org (`https://www.plainsite.org/tags/intellectual-ventures-shell-companies/`) to identifies subsidiaries of PAEs. Relying on these information, we reduce our sample to 3,580 entities.

documentation: “Up to grant of the European patent, transfers, licenses and other rights in respect of European patent applications are registered centrally in the European Patent Register in accordance with Rules 22 to 24 EPC. After grant of the European patent, a transfer is registered in the European Patent Register only during the opposition period or during opposition proceedings, in accordance with Rule 85 in conjunction with Rule 22 EPC”.[21]

This allows us to reconstruct the patent ownership histories during the entire granting process and thus to identify potential patent transfers within this period, which is crucial for analyzing the role of patent intermediaries such as PAEs.[22]

A change in applicant information registered in the EPR database reveals a potential patent transfer. However, as discussed by De Rassenfosse et al. (2017), not all communicated changes correspond to genuine transactions (just part of the registered changes should be considered as effective transactions).

EPR might register a patent transaction when in fact the event simply concerns a change in the firm name, given that names and addresses of the parties listed in the EPR database have not been harmonized or disambiguated. The very same applicant may thus have several customer identifiers, again leading to false positives in the analysis of patent transfers. To overcome this issue, we harmonize and standardize applicant names following a procedure described in Appendix A1.

Due to the relatively recent explosion of the PAE business, we restrict our analysis to EP patents filed during the period 1997-2012 (1,923,468 patents).[23] After applying the name cleaning and standardizing procedure described below and in Appendix A1, we individuate 314,490 unique patent

[21] `https://www.epo.org/law-practice/legal-texts/html/natlaw/en/ix/index.htm`.

[22] Precisely, we exploit information contained in the PATSTAT Register table 'REG107_PARTIES' – which provides data on applicants, inventors and legal representatives – to track changes of parties over time during the granting process. The types of parties are distinguished by the attribute 'TYPE'. For our purposes, we only consider applicants and inventors recorded, respectively, as 'A' and 'I'.

[23] We exclude patents filed after 2012 to ensure that there is sufficient time to observe both patent citations and transfers. We also exclude patents filed before 1997 since data on PAEs provided by Love et al. (2017) go back to 1997.

applicants listed in EP patent documents with filing year between 1997 and 2012. The total number of transferred patents is 244,437, representing 12.7% of the total sample. Within them, patents that are traded only once in their EPO life cycle come to 215,179 (88% of all transfer cases), those traded twice represent 10.3% of cases (25,242 patents) and more than twice 1.7% of cases (4,016 patents). For the purpose of our study, we focus our attention only on first transfers. Table 1 provides an overview of the phenomenon.

Table 1: Number of patent transfers registered at the EPR

Number of transfers	**Freq**	**Percent**	**Cum.**
0	1,679,031	87.29	87.29
1	215,179	11.19	98.48
2	25,242	1.31	99.79
3	3,398	0.18	99.97
4 or more	618	0.03	100.00
Total number of patents	1,923,468	100.00	

Years of patent filing: 1997-2012. Results obtained after cleaning and standardizing applicant's name and address. For the methodology description, see Appendix A1.

The PAE-EP database To identify EP applications assigned to PAEs, we perform a semantic matching procedure between entity names included in the aforementioned PAE list and the cleaned applicant names recorded in the EP-Register database.[24]

The matching procedure is a probabilistic matching which allows for a minimum amount of discrepancy between the applicant and PAE names to be matched. For the matching, we apply the RECLINK Stata algorithm (Blasnik, 2007).[25] This matching method leads to the identification of 12,598 EP patents in which at least one PAE appears as owner in the patent history,

[24]To perform the semantic matching we rely on cleaned-applicant names obtained as described in Appendix A1.

[25]We set the algorithm score at 0.95. This threshold has been chosen by visually comparing applicant names with PAEs names on a random sub-sample of 100 cases. For robustness checks we applied different thresholds (0.90 and 0.99): results do not change significantly and are available upon request by the authors.

representing 0.65% of the entire basket of EP patents filed from 1997 to 2012 at the EPO.[26] Table 2 shows the six categories we individuate at the EPR, according to first transfers and the type of applicant. Interestingly, this last descriptive evidence shows that 1,371 patents, directly applied at EPO by PAEs, have changed ownership during the granting process. The largest part of those patents have been transferred to PEs (1,131), while only 240 to other PAEs.

Table 2: Patent categories at EPR

Category	**Freq**	**Percent**	**Cum.**
Applied by PEs and never transferred	1,670,939	86.87	86.87
Applied by PEs and transferred to PEs	239,931	12.47	99.35
Applied by PAEs and never transferred	8,092	0.42	99.77
Applied by PEs and transferred to PAEs	3,135	0.16	99.93
Applied by PAEs and transferred to PEs	1,131	0.06	99.99
Applied by PAEs and transferred to PAEs	240	0.01	100.00
Total number of patents	1,923,468	100.00	

Years of patent filing: 1997-2012. Results obtained after cleaning and standardizing applicant name and address. For the methodology description, see Appendix A1. Only first transfers considered.

3.2 PAE activity at EPO: key figures

The industry PAEs essentially operate in ICT industries and, in general, in all the "complex" technologies (Kingston, 2001), in which a new product or process is composed of numerous separately patentable elements, leading to the fragmentation of the relevant IP ownership. This is confirmed by our data where, according to the 35-class OST patent classification (Schmoch, 2008),[27] the five most representative technological fields in which PAEs file or acquire patents are Digital Communication (21.9%), Computer Technology (11.4%), Telecommunications (10.8%), Audio and Visual Technology (8.9%) and Semiconductors (7.2%) (See Figure 1). For this reason, we decide to

[26]Individuated PAEs active at EPO during the period 1997-2012 come to 1047 (0.33% of the total number of registered applicants).

[27]`http://www.wipo.int/ipstats/en/statistics/technology_concordance.html`

exclude from the analysis patents in the low-tech sectors[28], even if we observe an increasing presence of PAEs in these domains during the last years[29]. Our final sample includes 7,633 PAE high-tech patents, representing 60.6% of all PAE patents.

Figure 1: PAE patent applications by technological field (1997-2010)

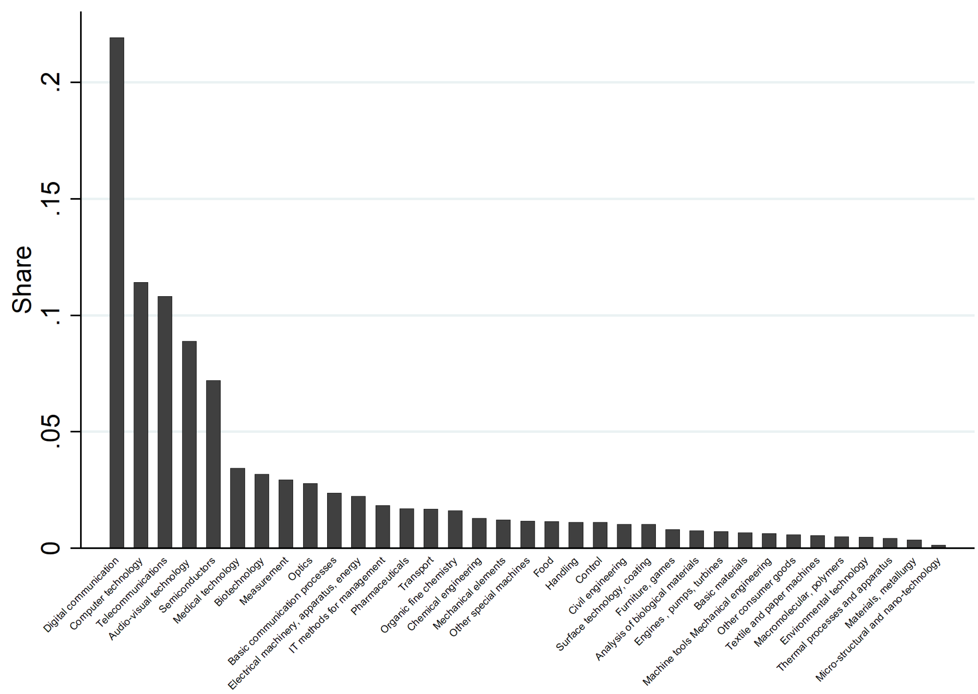

Notes: The figure plots the distribution of PAE-owned patents at EPO per main technological area (Schmoch, 2008)

[28]The definition of high-technology patents proposed by Eurostat uses specific subclasses of the International Patent Classification (IPC) as defined in the trilateral statistical report of the EPO, JPO and USPTO. The following (macro) technical fields are defined as high technology: Computer and automated business equipment; Microorganism and genetic engineering; Aviation; Communications technology; Semiconductors; Lasers. The list of sub-classes and their definition is provided by Eurostat at `http://ec.europa.eu/eurostat/cache/metadata/Annexes/pat_esms_an2.pdf`.

[29]If historically PAEs have been active mainly only in the high-tech sector, in the second half of the 2000s PAEs started a process of business differentiation and entered new markets, in part because low-tech industries have increased their use of computer-based technologies and in part because more low-tech companies started to sell patents to monetize their investment in R&D.

The way of entering the European patent market The way of entering the European patent market may be through the patent filing or through a patent acquisition. Focusing on high-tech patents with at least one PAE as owner, we find that the majority of them (5,612 – representing 73.5% of all PAE high-tech patents) are filed directly at the EPO by a PAE.[30] The rest (2,021 patents) has been acquired by PAEs from PEs. Restricting the focus on granted patents, we observe similar figures. Indeed, 65.7% of patents (or 1,743 patents) owned by PAEs have been directly filed at EPO by PAEs, while 34.3% (or 910) have been acquired by PAEs. Finally, by considering only transferred patents filed between 1997 and 2012, PAEs appear as owner of around 2.9% of them.. Table 3 summarizes those numbers.

Table 3: The PAEs' way of entering the market: filing and acquisition

	Applications	**(%)**	**Granted**	**(%)**
Filed by PAEs	5,612	*(73.5%)*	1,743	*(65.7%)*
Acquired by PAEs from PEs	2,021	*(26.5%)*	910	*(34.3%)*
Tot PAE patents	7,633		2,653	
Tot EPO patents	472,217		180,624	
Share of PAE patents	*(1.6%)*		*(1.5%)*	
Tot transferred EPO	70,341		31,544	
Share of transfers to PAEs	*(2.9%)*		*(2.9%)*	

Only High-tech EP patents are considered. Only first transfers considered. Years of filing: 1997-2012.

Age and granting process Patents transferred to PAEs are on average older than patents transferred to PEs (Table 4). The age of the invention at the time of the patent transfer, proxied by the years that elapse between the filing date and the transfer date, is on average 1.5 years higher for PAEs than for PEs. Furthermore, on average, PAE patents receive a grant later than PE patents (8.5 vs. 7.1 years after the filing date). These descriptive statistics suggest that, on average, PAEs and PEs follow different patent acquisition strategies.

[30]In most cases these are US patents that are likely to have been acquired by PAEs before being extended to the EPO.

Table 4: Average patent age at the first transfer and at the grant

	# of patents	Age (years)	Grant lag (years)
Acquired by PAEs	910	6.5	8.5
Acquired by PEs	30,317	5.0	7.1

Only granted high-tech transferred EP patents considered, originally applied by PEs. Years of filing: 1997-2012. Age is defined as the number of years that elapse between the filing date and the transfer date.

4 Empirical strategy and variables

To investigate the effect of PAE patent acquisition on innovation, we look at the number of citations received by patents transferred to PAEs, comparing them with both never-transferred patents and patents transferred to PEs.

Our sample comprises 178,564 unique high-tech granted patents filed at the EPO between 1997 and 2012. For each patent we collect information on its yearly number of citations received up to 2015,[31] building an unbalanced panel of 2,154,839 total observations.[32] We restrict our analysis to patents applied by PEs, excluding applications filed directly by PAEs which, in the majority of the cases, are foreign applications that have been acquired by PAEs and then extended at the EPO.[33] For these patents we do not observe in fact the moment in which they have been eventually transferred before the filing at the EPO.[34]

[31] Due to truncation issues, we collect citation data up to the year 2015.

[32] As described below, our sample reduces when we apply matching techniques.

[33] For example, 1,068 of the 1,743 granted high-tech EP patents that are firstly applied by PAEs and never transferred (61.3%) have at least one inventor resident in the US.

[34] However, also PEs may acquire patents in extra EU offices and then extend them at the EPO. For robustness, we replicate our baseline analysis described in Section 4.1 excluding patents invented outside the EU (those patents are indeed more likely to be extensions from other offices and thus transferred before being filed at the EPO). Results are consistent with the ones discussed in Section 5.1 and reported in Table 7, and are available upon request by the authors.

4.1 Forward citations as an indicator of patent technological quality and exploitation

We assess the impact of PAEs on innovation by firstly looking at the pattern of patent citations the focal (transferred) patent receives. We consider patent citations[35] both as an indicator of patent exploitation and a measure of technological quality. We argue that the number of forward citations is an indicator of the fact that the patented technology is somehow used by innovating and producing companies (Trajtenberg, 1990), whether they are patent holders (or licensees) or other companies performing R&D activities, or both. Citations are reported in the patent document, provide a legal delimitation of the property right scope, and have been used in the literature to track knowledge flows (Jaffe et al., 1993; Jaffe and Trajtenberg, 1999; Maurseth and Verspagen, 2002; Bottazzi and Peri, 2003; Bacchiocchi and Montobbio, 2010; Montobbio and Sterzi, 2011).[36] Since citations show the degree of novelty and the inventive steps of patent claims, they identify the antecedents upon which the invention stands. In this respect, a citation from patent A to patent B indicates that part of the knowledge protected by patent B is also used in the technology protected by patent A. Controlling for the age and the domain, patents that stop being cited indicate that the protected technologies are likely to be no longer used in further inventions. Conversely, a high number of citations received indicates that the patented invention is of high technological quality (Trajtenberg, 1990; Fischer and Leidinger, 2014).

In this subsection, we firstly examine the distribution of forward citations received by, respectively, patents transferred to PAEs, patents transferred to PEs and patents that never change ownership during their life cycle, focusing on the citation age profile – i.e., the average citation rate conditional on the age (since the filing at the EPO) of the cited patent. When comparing the citation age profile of these three categories we do not use any information about the timing of the transfer (which may occur at any point during the life of a transferred patent). Thus, we simply focus on the question of

[35] We correct forward citations for DOCDB patent families to avoid multiple counting (Martínez, 2011).

[36] Griliches (1998) and Breschi et al. (2005) provide path-breaking and renowned surveys on the topic.

whether transferred patents (and, particularly, patents transferred to PAEs) are especially significant.

As a second step of the descriptive analysis, we restrict our focus to only transferred patents, with the aim of comparing the average age profile (before and after the transfer) of patents transferred to PEs and patents transferred to PAEs.

Figure 2 illustrates our initial findings. While patents transferred to PEs and never transferred patents show a similar distribution, patents that at some point of their life are transferred to PAEs show, on average, a high number of citations in their earlier phase of life, followed by as much decline from the third year on.

In Figure 3 we restrict our focus to transferred patents and we compare the citation age profile before and after the transfer for two groups of interest, namely PE-transferred and PAE-transferred patents. We follow the transferred patents in a two 5-years windows, pre- and post-transfer. It is worth noticing that patents transferred to PEs show a steady citation rate before the transfer year, while they experience a drop after. Our main group of interest (patents transferred to PAEs), instead, presents a different citation profile than the former. In the pre-transfer period, this group indeed receives on average more citations than the other. However, we observe the citation profile of those patents falling below the profile of patents transferred to PEs in the after-transfer period. Moreover, the drop in citations for patents transferred to PAEs starts around two years before the transfer occurs.

4.2 A triple differences approach (DDD)

To study the impact of PAEs on innovation dynamics we firstly rely on a triple differences (DDD) research design in a panel data framework with patents that experience a change of ownership as the treated group and patents that are never transferred as the control group. We further split the treated group of patents into two groups: (1) patents transferred to PEs (PE); and (2) patents transferred to PAEs (PAE). Our patent-level DDD setup accounts for common macroeconomic trends and observable technological characteristics. This specification allows us to examine the difference

Figure 2: Age profile of citations since the patent filing (all categories)

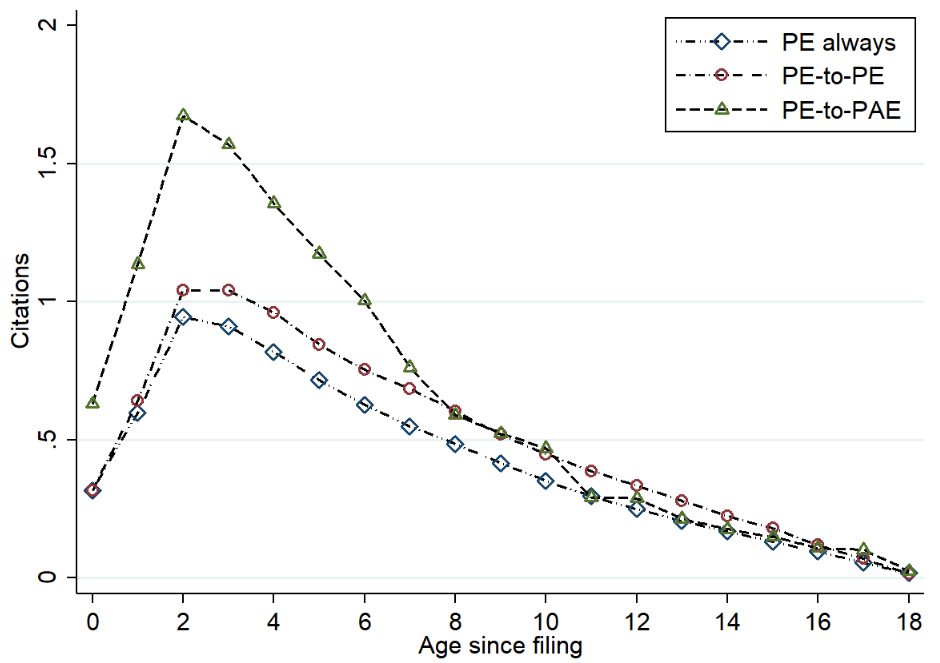

Notes: The figure draws the age profile of citations since the patent filing by category (never sold patents, patents transferred to PEs and patents transferred to PAEs). The category "PE Always" corresponds to patent applications that are never transferred during their life; the category "PE-to-PE" corresponds to patent applications filed by PEs and that are transferred at some point during the patent life to other PEs; the category "PE-to-PAE" corresponds to patent applications filed by PEs and transferred during the patent life to PAEs.

Figure 3: Age profile of citations before and after the patent transfer (transferred patents)

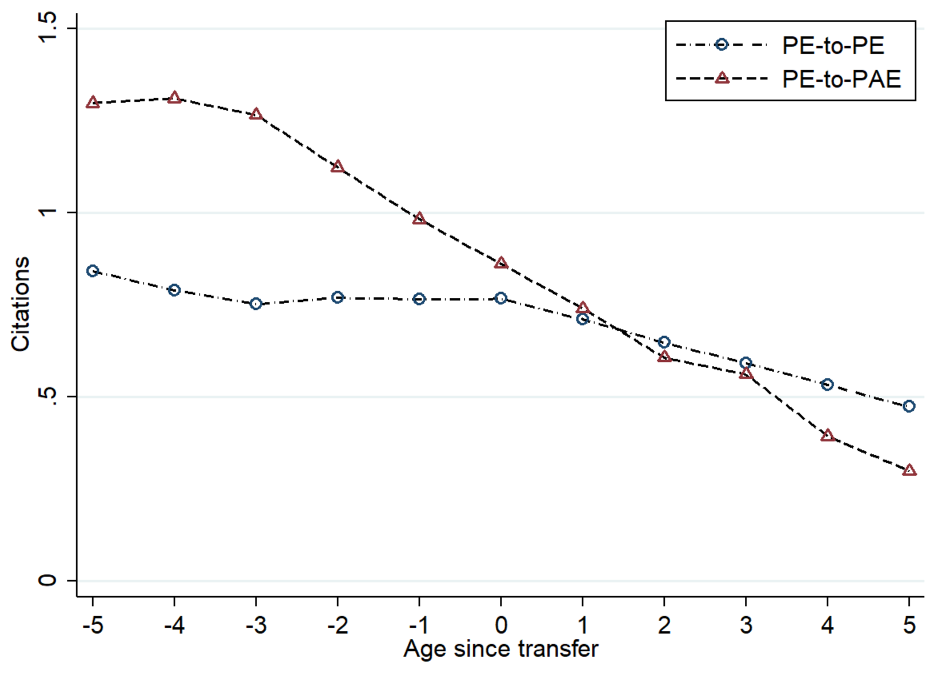

Notes: The figure draws the age profile of citations since transfer by category (patents transferred to PEs and patents transferred to PAEs).

between the change in innovation diffusion by patents acquired by PAEs and the corresponding change by patents acquired by PEs.

To estimate the impact of PAEs on innovation dynamics, we estimate the following empirical model to predict the yearly number of citations received by the patent i during its life.

$$
\begin{aligned}
Cit_{it} \;=\; & \alpha_0 + \alpha_1 PE_i + \alpha_2 PAE_i + \alpha_3 TRADED_{it} + \alpha_4 TRADED_{it} * PAE_i + \\
& + \sum \beta_t Age_t + \sum \gamma_l FilingYear_l + \mathbf{X_i'}\delta + f_i + \varepsilon_{it} \qquad (1)
\end{aligned}
$$

where Cit_{it} is the number of citations received by patent i in the year t. We take the log (plus one) of the number of citations to have the dependent variable more closely distributed to normality. PE_i is a dummy variable to indicate patents that are transferred to PEs, while PAE_i is a dummy variable for patents transferred to a PAE. The reference group is composed of patents applied by PEs and never traded (the group labeled "*PE always*" in Figure 2). $TRADED_{it}$ is an indicator of the post-traded event related to the first transfer: it is a dummy variable that identifies the change of ownership for each patent such that it is always zero for patents that are never transferred, while it takes the value one for transferred patents from the year of the transfer and in subsequent years. On one side, a positive sign of the dummy $TRADED_{it}$ might indicate that patent transactions facilitate better matches between technology suppliers and users. On the other side, a negative sign might indicate that patents are used and acquired mainly for strategic reasons (Hall and Ziedonis, 2001; Blind et al., 2009; Noel and Schankerman, 2013), as it happens especially in the case of complex technologies (Bessen, 2003; Orsenigo and Sterzi, 2010).

To control for the technology life cycle, we include dummies related to the patent priority year ($FilingYear$) and dummies (Age) for each year since the patent's priority filing (which is normalized to zero). Note that while dummies Age vary within patent, dummies $FilingYear$ do not. $\mathbf{X_i'}$ is a vector of patent fixed characteristics that are potentially associated with patent forward citations. Their inclusion may improve the accuracy of the DDD estimate. Among these controls, we include a set of technological field dummies, the inventor team size, the level of patent originality,[37] the dum-

[37]Patent originality is calculated according to Squicciarini et al. (2013). Quoting

mies for the inventor's country of residence, the number of patent claims, and dummies for patents applied by individuals and those applied simultaneously by more than one applicant; ε_{it} is the error term.

The description of the variables used in the empirical analysis and their sources are presented in Table 4.2. Summary statistics are presented in Table 6. DDD results are presented in Section 5.

Parameters α_1 and α_2 measure the difference in the average number of forward citations computed in the period before the transfer between the reference group (never transferred patents) and, respectively, the group of patents transferred to PEs (PE_i) and the group of patents transferred to PAEs (PAE_i). Positive signs for these parameters indicate that patents that will be transferred during their life are on average of higher quality than patents never transferred. In particular, a positive sign of α_2 indicates that PAEs "cherry pick" high-quality patents.

Our main interest focuses on parameters α_3 and α_4. The parameter α_3 identifies the effect of market transactions on citation dynamics, when the buyer is a PE. Two main forces drive the sign of this coefficient. On the one side, the (secondary) patent market is likely to facilitate the match between buyers and sellers, so that the patent transaction promotes innovation (positive effect). On the other side, if the patent is acquired for strategic reasons, its transfer will be detrimental to its further usage (negative effect). The parameter α_4 is the difference-in-difference-in-differences estimator and identifies the impact of PAEs on citations. A positive sign of α_4 indicates that transferring a patent to a PAE, rather than to a PE, increases the chance that the technology protected by the patent will be subsequently used and exploited by innovating firms, thus cited more frequently (*"market-makers"* view). On the contrary, a negative sign indicates that patents acquired by

the authors, *"Patent originality refers to the breadth of the technology fields on which a patent relies. The patent originality measure, first proposed by Trajtenberg et al. (1997), operationalizes this concept of knowledge diversification and its importance for innovation: inventions relying on a large number of diverse knowledge sources are supposed to lead to original results (i.e. on patents belonging to a wide array of technology fields)"* [pag. 49]. Building on Hall et al. (2001), they define the originality indicator as: $Originality_p = 1 - \sum_j^{n_p} s_{pj}^2$. where s_{pj} is the percentage of citations made by patent p to patent class j out of the n_p IPC 4-digit patent codes contained in the patents cited by patent p. Citation measures are built on EPO patents and account for patent equivalents.

Table 5: Variables description and sources

Variable	Definition	Source
CITATIONS	Number of EP forward citations corrected for DOCDB patent families	CRIOS-PATSTAT
PE	Patent acquired by a PE (dummy)	European Patent Register
PAE	Patent acquired by a PAE (dummy)	European Patent Register
TRADED	Indicator of the post-traded event (=1 from the year of the transfer)	European Patent Register
AGE	Number of years elapsed since the patent filing	European Patent Register
COAPPLICANT	Patent applied by two or more applicants (dummy)	European Patent Register
INDIVIDUAL	Patent applied by an individual (dummy)	European Patent Register
TEAM SIZE	Number of inventors	European Patent Register
ORIGINALITY	Patent originality index	OECD Patent Quality Indicators
CLAIMS	Number of claims in a patent document	OECD Patent Quality Indicators
PATENT STOCK	Applicant stock of patents at the time of the patent filing	European Patent Register
POOL	Patent transferred with at least other 24 patents (dummy)	European Patent Register
COUNTRY	Patent inventor's country of residence (dummy)	European Patent Register
TECHNOLOGY	OST7 technological domain (dummy)	CRIOS-PATSTAT

CRIOS-PATSTAT (Coffano and Tarasconi, 2014) ; OECD Patent Quality Indicators (Squicciarini et al., 2013).

Table 6: Summary Statistics

	Mean	**St. Dev.**	**Min.**	**Max.**
CITATIONS	0.6	1.63	0	102
PE	0.2	0.39	0	1
PAE	0.01	0.07	0	1
TRADED	0.1	0.32	0	1
AGE	6.3	4.50	0	18
COAPPLICANT	0.05	0.22	0	1
INDIVIDUAL	0.02	0.12	0	1
TEAM SIZE	2.7	1.84	1	29
ORIGINALITY	0.7	0.22	0	1.0
CLAIMS	13.8	9.63	1	182
PATENT STOCK (thousands)	0.23	0.82	0	10.29
POOL	0.1	0.30	0	1
Observations	2,160,024			

Patent filing year between 1997 and 2012. For the variables description and sources, see Table 4.2.

PAEs start receiving fewer citations after the transfer as compared to patents sold to PEs, suggesting that PAEs do not facilitate cumulative innovation, but rather stand in its way (*"market-breakers"* view). Finally, the combination of α_3 and α_4 identifies the effect of market transactions on the citations path when the patent is acquired by a PAE.

4.3 Propensity score matching and conditional DDD

One might reasonably question that the fact that the decision to transfer a patent is not exogenous. Exploiting the longitudinal dimension of our data guarantees that relevant issues related to unobservable factors are taken into account. By adding patent patent fixed effects to model (1) we control for all time-invariant unobservable patent characteristics. However, a bias due to observable variables is likely to still remain. For example, patent characteristics such as the age of the patent, the number of citations received by a given age, and the patent generality may influence the probability that a patent is transferred (Serrano, 2010).

We may expect, for example, that companies to target patents in high-growth technological sub-domains, resulting in an increasing trend in the citations path after the transfer occurs and implying a positive bias in the coefficient for the dummy $TRADED_{it}$.

In presence of potential biases due to selection into treatment, the DDD model may produce non-consistent estimates, even when it controls for observed variables that might influence both the outcome and the treatment. To partially overcome biases due to observable factors, we apply matching methods. Matching methods seek to replicate a randomized experiment in which the matched and the control patents do not differ systematically from each other on observable variables. Consequently, we match patents that are transferred (either to PAEs or to PEs) and non-transferred patents on an index, the propensity score, of several characteristics affecting the likelihood of a transfer occurring. We assume that all variables relevant to the probabilty of observing a transfer are observed and included in the model so that we may construct an unbiased counterfactual of non-transferred patents for the group of traded patents (conditional independence assumption).

More precisely, among these observable characteristics, we include: the patent filing year; the average number of citations received in the 4-year time window elapsing from the filing[38] and the level of patent originality, as proxies for the patent technological quality; the number of backward citations; the technological sub-field in which the invention belongs to (accounting for intrinsic technological fixed effects); the number of patent claims (as a proxy for the patent scope); the size of the inventor team; the nature of the applicant (individual vs. company); whether the patent is co-applied; the size of the first applicant (proxied by the applicant's stock of patents);[39] and the inventor's country of residence. Since we look at the first four years after the filing to both count the number of forward citations and measure the patent technological quality, we drop patents that have been transferred within this time window from the analysis.

The propensity score is then calculated from the fitted values of a probit model where the dependent variable is the probability of a patent transfer to occur. We adopt the nearest-neighbor algorithm, using the information from up to five neighbors and setting a "caliper" threshold to 0.02. As Caliendo and Kopeinig (2008) illustrate, the choice of the algorithm to use is a matter of a trade-off between bias and efficiency. Using up to 5 control units to proxy for the counterfactual situation allows us to gain efficiency in the estimation, while the caliper threshold, which imposes a tolerance level on the maximum propensity score distance, reduces potential bias, avoiding bad matches.[40]

Through matching techniques we restrict the analysis to treated and control patents that are on average observationally almost identical on a set of fixed characteristics. The final restricted sample, composed of transferred-

[38]The choice of considering four years for citations is due, on the one hand, to the fact that patents receive the majority of citations in the first four years from the filing and, on the other hand, to the fact that the first transfer occurs, on average, after four years when PEs are buyers. For robustness we also count citations only up to the second year after the filing: results are very similar to those presented in Table 9 and available upon request by the authors.

[39]The patent stock is calculated applying the Perpetual Inventory Method, with a 15% annual rate of obsolescence.

[40]The selected caliper value is very conservative and corresponds approximately to 0.1 times the standard deviation of the propensity scores recovered with the probit regression. For robustness, we re-run our estimates using diverse tresholds (up to 0.25 times the propensity scores recovered with the probit regression). Results remain almost identical to the ones presented in Section 5.3 and are available upon request by the authors.

patents and similar patents that are non transferred, is then used for estimating model (1), so that we follow a conditional difference-in-difference-in-differences (CDDD) strategy.

4.4 A finer analysis of transferred patents

While the analyses described so far lead us to interpret the role of PAEs from a comprehensive perspective, we acknowledge that they come at the cost of not entirely solving endogeneity issues. The patent transfer is indeed an endogenous event since we cannot properly control for entity strategies. Even if we are close to replicating a hypothetical experiment by both performing matching techniques and exploiting the longitudinal nature of our data, an intrinsic source of bias remains.

In particular, as for different strategies followed by PEs and PAEs in patent purchases, while PEs might mainly target patents that protect technologies that are strategic for their R&D activities, we argue that PAEs are instead more likely to target patents just focusing on their usefulness for suing (producing) companies for infringement. The likelihood of receiving citations for PAEs' patents – as an indicator of technological patent quality – may thus, in this case, be systematically different than for PEs' patents (more likely lower). If this argument leads us to assume that PAEs target patents with scarce technological content, other arguments may tell an opposite story. For example, PEs are involved several times in transactions in which patents are just complementary assets, not necessarily the core of the deal (i.e. M&As). PAEs' structure and strategy are instead essentially built to either acquire patents to monetize them or inherit patents from unsuccessful operating companies (Shapiro and Scott-Morton, 2014; Scott Morton and Shapiro, 2016). Since patents are almost the only asset PAEs have, it is reasonable to assume that they may be more accurate than PEs in building up their patent portfolios. As a result, the average qualitative level of patents acquired by PAEs may be systematically higher than the level of those acquired by PEs.

To more deeply investigate the effect of patent transfers to PAEs on their forward citation path, we restrict our analysis to only transferred patents. Starting from the descriptive evidence provided in Figure 3, we perform a difference-in-differences analysis considering, as treated patents, the patents that have been transferred to PAEs and, as control patents, the ones that have been transferred to PEs.

For the similar reasons stressed in the previous subsection and to minimize possible biases due to observable factors, we apply propensity score matching techniques on several fixed, pre-treatment patent characteristics when selecting the patents constituting our control group. Precisely, we add two covariates to those used in the previous analysis: the year of the transfer and a dummy variable indicating whether the transferred patent was (presumably) part of a patent portfolio acquisition ($POOL_i$).[41] In particular, controlling for the year of the transfer (along with the age and the technological field of the patent) allows us to compare patents that are following a similar path at the time of the transfer.

As before, since we look at the first four years after the filing to build our main proxy of patent technological quality, we drop patents that have been transferred in the first four years after the filing.

The propensity score is then calculated from the fitted values of a probit model where the dependent variable is the probability that a patent is transferred to a PAE. We adopt the nearest-neighbor algorithm, using the information from up to five neighbors and setting a caliper threshold to 0.02.

Since selection on unobservables may represents a relevant concern and might bias the estimation, we maintain the longitudinal structure of the data and we follow the fixed effects conditional difference-in-difference (CDD) strategy. Our model takes the following form:

$$Cit_{it} = \alpha_i + \beta_1 TRADED_{it} + \beta_2 PAE_i * TRADED_{it} + \sum \gamma_t Age_t + \varepsilon_{it} \quad (2)$$

where Cit_{it} is the number of citations received by patent i at time t; PAE_i is an indicator for whether a patent has been transferred to a PAE; $TRADED_{it}$ indicates whether the observation belongs to post-transfer periods. As before, we include dummies (Age_{it}) for each year since the patent's priority filing (which is normalized to zero); α_i is a patent fixed effect (which absorbs any time-invariant characteristic including the "main effect"

[41] We consider a transfer from a seller s to a buyer b in year t as a transfer occurring in a pool if, in the same year, the same buyer b acquires at least 25 patents. Our estimates are robust to different thresholds and are available upon request by the authors.

of PAE_i); and ε_{it} is the error term. The parameter β_2 represents the average causal effect of a transfer to a PAE, with respect to a PE, on patent citations.

5 Results

In this section we present the results of the empirical approaches proposed above. We begin in Section 5.1 with a baseline evaluation of the impact of a PAE patent acquisition on innovation using the patent-level DDD research design described in Section 4.2. In Section 5.2 we test the robustness of the baseline results by refining our measures based on patent citations. We then question the exogeneity of the patent transfer and present the results from the CDDD fixed effect models in Section 5.3. Finally, in Section 5.4 we restrict our focus to transferred patents and we more deeply investigate differences between citations received by patents transferred to PAEs and citations received by patents transferred to PEs.

5.1 Baseline Results

With respect to Equation 1, we take the logarithm transformation of the dependent variable and we estimate OLS models.[42] We cluster standard errors at the patent level to control for possible serial correlations (Bertrand et al., 2004). Table 7 presents the estimation results. Different specifications refer to the inclusion in the specification of different controls. To interpret the magnitude of the coefficients, we refer to model (4), which contains the full set of control variables.

The coefficient related to the dummy PE_i is significant and positive. Holding everything else constant, PEs acquire patents that are above the average in terms of citations received (precisely, before the time of the transfer, those patents receive 1.5% more citations per year than never transferred patents). Interestingly, patents transferred to PAEs (dummy PAE_i) receive on average 11% more citations, before the transfer, than patents never sold

[42] In the logarithm transformation we add one to all values. As patent counts take only non-negative integer values, we further estimate count models which give similar results and are presented in Table 18 (Appendix A3).

in the patent market, meaning that patents acquired by PAEs are on average of higher quality than those never transferred.[43] This result, in line with Fischer and Henkel (2012) and Leiponena and Delcamp (2018), is in conflict with the common feeling that PAEs' patent portfolios are mainly constituted of sparse and low-quality technologies.

The dummy $TRADED_{it}$ is an indicator of the post-traded event and, when it is not interacted with the dummy PAE_i, it refers to patents sold to practicing entities. The associated parameter (α_3 in the Equation 1) is not statistically significant, meaning that the transfer to a PE does not affect the rate of citations the patent receives.

The interaction term $TRADED_{it} * PAE_i$ identifies the additional effect of the transfer when the trade involves a PAE. In model (4), the coefficient is -0.098, meaning that, after the transfer, patents transferred to PAEs receive around 9.8% fewer citations than patents transferred to PEs. Since α_3 is not significant, the net effect of the transfer to a PAE on the patent citation rate is thus negative, implying a reduction of 9.8% in the number of citations per year in the post-transfer period; this implies that, after the transfer, patents transferred to PAEs and never-transferred patents are no more statistically different in quality. These results seem to corroborate the idea of PAEs as *"market-breakers"*, discussed in Section 2.2. However, given the endogeneity concerns, the negative sign of PAE_i may indicate either that (a) PAEs create an obstacle to innovation by discouraging producing companies from entering and investing in fields related to the transferred patent or that (b) PAEs specifically target patents which are already in the declining phase of their life cycle.

5.2 Strategic citations and the "in house" effect

On the one hand, results presented in Section 5.1 indicate that PAEs acquire patents that are, on average, of high technological quality. On the other hand, they show that those patents receive fewer citations after the transfer, confirming the "*Market-breakers*" view presented in Section 2.2.

[43] Coefficients about the dummies PE and PAE are statistically different, meaning that the average quality of patents transferred to PAEs is also higher than the average quality of patents transferred to PEs.

Table 7: Baseline models

	(1)	(2)	(3)	(4)
	Cit (LN)	**Cit (LN)**	**Cit (LN)**	**Cit (LN)**
PE	0.054***	0.052***	0.014***	0.015***
	(0.0031)	(0.0031)	(0.0030)	(0.0031)
PAE	0.13***	0.18***	0.12***	0.11***
	(0.014)	(0.018)	(0.018)	(0.018)
TRADED	-0.041***	-0.037***	-0.0014	-0.0015
	(0.0030)	(0.0031)	(0.0030)	(0.0030)
TRADED*PAE		-0.10***	-0.098***	-0.098***
		(0.017)	(0.016)	(0.016)
TEAM SIZE (LN)				0.059***
				(0.0019)
ORIGINALITY				0.084***
				(0.0033)
CLAIMS (LN)				0.083***
				(0.0014)
COAPPLICANT				-0.0032
				(0.0038)
INDIVIDUAL				-0.012*
				(0.0064)
PATENT STOCK (LN)				0.0020***
				(0.00028)
Age FE	Yes	Yes	Yes	Yes
Filing Year FE	No	No	Yes	Yes
Technology FE	No	No	Yes	Yes
Country FE	No	No	Yes	Yes
Observations	2,154,839	2,154,839	2,154,839	2,154,839
Number of patents	178,564	178,564	178,564	178,564
Adjusted R^2	0.049	0.049	0.088	0.101
F	118.8	93.5	95.4	202.0

Column (1) reports our most parsimonious specification, without our interaction of interest and with only patent age dummies as covariates. In column (2) we add our interaction of interest. In column (3) we also control for a series of dummies: patent filing year, inventor's country of residence and technological domain. Column (4) is our preferred specification in which we add the full set of covariates. In column (5) we include patent fixed effects (and exclude time invariant controls). Clustered Standard errors at the patent level are in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

To assess the robustness of these results, we first exclude the number of citations added by the applicant from the total number of forward citations. Indeed, one might think that actors involved in R&D projects in fields related to those in which PAEs are active may strategically decide not to cite patents owned by PAEs, if they perceive an augmented risk of being sued. If this is the case, we would over-estimate the overall negative effect of PAEs' patent acquisitions on follow-on innovation activities: the patent purchase by a PAE would impact only the citation paths without really reducing innovation. In order to discard this possibility, we consider only citations added by the patent examiner. The results proposed in Table 8 confirm the general ones, revealing that this source of bias is only marginally present. The coefficient of the interaction term $TRADED_{it} * PAE_i$ is still significant and negative, although it decreases from -9.8% (Table 7, Column 4) to -9.3% (Table 8, Column 4).

As a further robustness test, we also exclude self-citations at the applicant level (those where citing and cited applicants are the same) from the count of forward citations. Results are reported in Appendix A3 (Table 19) and confirm the extant evidence.[44]

The evidence provided by this first part of the analysis suggests that PAEs target patents revealing high-quality technological content.[45] This results conflicts with the idea that patents owned by PAEs are, on average, of lower technological value and weaker than patents owned by producing companies. However, once acquired, those patents experience a strong decline in their citation path, irrespective of the way we count forward citations. This evidence tends to strengthen the *"Market-breakers"* view stated in Section 2.2, posing PAEs in a negative position in terms of impact on innovation.

---

[44] Importantly for our analysis, when we exclude self-citations from the count, we mainly capture external knowledge spillovers and we discard the in-house development based on the acquired patent. In this sense, the flow of citations is rather indicative of external use of the technology protected by the patent. Note that PAEs are, by definition, non-practicing entities that do not acquire patents for internal use. Conversely, net of strategic operations, PEs do acquire patents to build on further inventions. By excluding self-citations from the count of forward citations we look more precisely at the impact of trade on knowledge diffusion.

[45] Coefficients about PE and PAE are statistically different, meaning that PAEs target patents even qualitatively better, on average, than those acquired by PEs.

Table 8: Baseline models (Exclusion of Citations added by the applicant)

	(1)	(2)	(3)	(4)
	Cit (LN)	**Cit (LN)**	**Cit (LN)**	**Cit (LN)**
PE	0.053***	0.051***	0.013***	0.014***
	(0.0030)	(0.0030)	(0.0030)	(0.0030)
PAE	0.13***	0.18***	0.11***	0.11***
	(0.013)	(0.018)	(0.018)	(0.018)
TRADED	-0.041***	-0.038***	-0.0017	-0.0018
	(0.0029)	(0.0030)	(0.0029)	(0.0029)
TRADED*PAE		-0.100***	-0.094***	-0.093***
		(0.017)	(0.016)	(0.016)
TEAM SIZE (LN)				0.058***
				(0.0019)
ORIGINALITY				0.080***
				(0.0032)
CLAIMS (LN)				0.081***
				(0.0014)
COAPPLICANT				-0.0038
				(0.0036)
INDIVIDUAL				-0.012*
				(0.0061)
PATENT STOCK (LN)				0.0018***
				(0.00027)
Age FE	Yes	Yes	Yes	Yes
Filing Year FE	No	No	Yes	Yes
Technology FE	No	No	Yes	Yes
Country FE	No	No	Yes	Yes
Observations	2,154,839	2,154,839	2,154,839	2,154,839
Number of patents	178,564	178,564	178,564	178,564
Adjusted R^2	0.050	0.050	0.090	0.103
F	119.8	93.6	99.6	205.2

All the models use the count of applicant-excluded forward citations as the dependent variable. Column (1) reports our most parsimonious estimation without our interaction of interest, with only patent age fixed effects included. In column (2) we add our interaction of interest. In column (3) we also control for a series of dummies: patent filing year, inventor's country of residence and technological domain. Column (4) is our preferred specification in which we add the full set of covariates. In column (5) we include patent fixed effects (and exclude time invariant controls). Clustered Standard errors at the patent level are in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

5.3 PSM and CDDD

As highlighted in Sections 4.3 and 4.4, one might question that the fact that a patent is traded is not exogenous. To partially overcome this source of bias further, we apply matching methods, seeking to replicate a randomized experiment in which the matched and the control patents do not differ systematically from each other on observable variables. More precisely, we match patents that are transferred (either to PAEs or to PEs) and non-transferred patents on an index, the propensity score, of several characteristics affecting the likelihood of a transfer occurring (Serrano, 2010). The list of variables selected to perform the matching concerns a comprehensive set of patent, applicant and inventor characteristics (see Section 4.3). The tests performed on the quality of the matching reveals that the adopted procedure successfully corrects for the selection on observable factors (Appendix A2.1 is dedicated to an in-depth analysis of the tests performed for assessing the quality of the matching). As described in Section 4.3, in this part of the analysis we do not consider transfers occurred in the first four years from the filing. Our group of transferred patents reduces to 19,311 (of them, 749 have been transferred to PAEs).[46]

Once the propensity scores are calculated and the quality of the adopted matching procedure assessed, we present the results of the CDDD estimation in Table 9. We replicate the strategy proposed in Section 4.2 over the reduced sample (81,179 patents) resulting from the PSM. We thus again estimate Equation 1, including age and patent fixed effects.[47] Our dependent variables (and all the variables related to patent citations that we use for implementing the relative matching) are, alternatively, the raw count of citations (Column 1), the count of citations with the exclusion of those

[46]Within the 31,544 high-tech tranferred EP patents filed between 1997 and 2012 ("treated" group), we do not find a suitable control for 12,233 cases (38.8%). Of them, 12,072 are patents transferred to PEs (39.4% of the total number of patents transferred to PAEs), while 161 are patents transferred to PAEs (17.7% of the total number of patents transferred to PAEs). The control group (never transferred patents) is formed by 61,868 patents.

[47]Results based on the other specifications do not show relevant differences and are available upon request by the authors.

added by the patent applicant (Column 2), and the count of citations with the exclusion of self-citations (Column 3).

Results, reported in Table 9, confirm the main findings highlighted in Section 5.1 and in Section 5.2. The main difference from the baseline results is the role played by the dummy $TRADED_{it}$. Once patents that are transferred (either to PAEs or to PEs) have been matched to non-transferred patents on observable fixed characteristics and once we control for patent fixed effects, we do indeed find that, after the transfer, patents receive fewer citations. This result may be explained by the increasing incidence of strategic patent acquisitions in the ICT domain (Hall and Ziedonis, 2001; Blind et al., 2009; Noel and Schankerman, 2013; Torrisi et al., 2016). Strong technological complementary and standardization, typical of the high-tech sector, lead to a mutual hold-up among innovators and to the fragmentation of the relevant IP ownership (Orsenigo and Sterzi, 2010), so that the exploitation of cross-licensing agreements and the ability of both avoiding hold-up problems and attracting venture capital funding are often the main reasons for patent acquisitions (Hall and Ziedonis, 2001). Thus, PEs often acquire patents that they "do not use". This "non-use" is mainly linked to the acquisition of preemptive patents, which may serve for blocking competitors or for ensuring the freedom to operate (Walsh et al., 2016).

Moreover, the (additional) effect of a patent transfer to a PAE on the follow-on use of the protected technology is negative and significant in all the specifications. Precisely, according to the estimates reported in Column 1 (where the dependent variable is the log-transformed raw number of forward citations), the transfer of a patent to a PAE reduces the yearly number of forward citations it will receive by 11% compared to patents transferred to practicing entities. Comparing this result with the ones from the baseline estimations (Table 7, Column 5), we find that the net effect is even more negative, going from -11% to -11.8%.[48]

[48]The same evidence appears when comparing the estimates from Column 2 with estimates from Column 5 in Table 8, and when comparing the estimates from Column 3 with estimates from Column 5 in Table 19.

Table 9: CDDD models (fixed effects)

	(1)	(2)	(3)
	Cit raw (LN)	**Cit No Appl (LN)**	**Cit No Self (LN)**
TRADED	-0.0080***	-0.0065**	-0.0071**
	(0.0028)	(0.0028)	(0.0028)
TRADED*PAE	-0.11***	-0.11***	-0.11***
	(0.016)	(0.015)	(0.015)
Age FE	Yes	Yes	Yes
Patent FE	Yes	Yes	Yes
Observations	1,088,931	1,088,931	1,088,931
Number of patents	81,179	81,179	81,179
Adjusted R^2	0.405	0.400	0.399
F	32.5	31.5	32.5

Model (1) uses the raw count of forward citations as the dependent variable. Model (2) uses the count of applicant-excluded forward citations as the dependent variable. Finally, Model (3) uses the count of self-citation-excluded forward citations as the dependent variable. All the models include patent and age fixed effects. Clustered Standard errors at the patent level are in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

5.4 A finer analysis of transferred patents

In this subsection, we more deeply investigate the effect of patent transfers to PAEs on their forward citation path. To do so, we restrict the sample to only high-tech granted EP patents, applied by PEs during the period 1997-2012, that have been transferred at least once after their first filing at EPO. We then match patents transferred to PAEs with a control group of patents transferred to PEs. Precisely, as discussed in Section 4.4, we perform a propensity score matching on several pre-transfer fixed patent characteristics, with the aim of narrowing down our sample to patents that share similar characteristics at the time of the transfer.

This procedure leads to a reduction in the number of observations. Indeed, we match 866 treated patents with 3,500 controls.[49] Appendix A2.2 reports the tests performed on the quality of the matching procedure applied.

After performing the matching, we estimate the impact of a patent transfer to a PAE on the patent forward citation pattern in a difference-in-difference framework. Accordingly, we estimate equation 2. Results are reported in Table 10. Columns 1 shows the estimated coefficients of our main interaction of interest when the dependent variable is the log-transformed number of total citations received by the focal patent. In the second and third columns, our dependent variable is, respectively, the log-transformed number of citations exclusively added by the patent examiner (excluding cites added by the patent applicant) and the log-transformed number of citations net of self-citations. All the estimations include patent and age fixed effects.

Results are not sensitive to the way we compute the number of forward citations: patents do receive fewer citations after the transfer and we do not observe a statistical difference according to the type of buyer. Taken as a whole, the evidence shows that, when we compare patents that are similar at the time of the transfer, the impact on the citation rate of a PAE's entry is not statistically different from the impact of a PE's entry. Even if negative, the coefficient of our interaction of interest, β_2, is always non significant. This result suggests that the negative impact of PAEs on innovation found

[49] For the reasons discussed in Section 4.3 and 4.4, we do not consider here patents transferred during the four years since the filing. Our sample of treated patents thus reduces from 910 to 888. Of them, 22 do not match with any of the control patents.

in the previous analysis is driven, at least partially, by the fact that PAEs and PEs do not target similar patents: PAEs mainly acquire patents that are already in the declining phase of their technology life cycle. This last result goes in the direction of, possibly, attenuating the statement we propose with the *"market-breakers"* view in Section 2.2. Although our results indicate that PAEs do not behave as intermediaries, we may indeed not conclude that they create an additional obstacle to innovation by discouraging producing companies from entering and investing in fields related to the transferred patent: in fact, we do observe that the negative impact of the transfer on the citation path of transferred patents is not uniquely attached to PAEs but, interestingly, verifies also when PEs are involved in patent purchases.

Table 10: Transferred patents results (CDD models, fixed effects)

	(1)	(2)	(3)
	Cit raw (LN)	**Cit No Appl (LN)**	**Cit No Self (LN)**
TRADED	-0.046***	-0.044***	-0.044***
	(0.012)	(0.012)	(0.012)
TRADED*PAE	-0.014	-0.012	-0.014
	(0.017)	(0.016)	(0.016)
Age FE	Yes	Yes	Yes
Patent FE	Yes	Yes	Yes
Observations	57,162	57,162	57,162
Number of patents	4,366	4,366	4,366
Adjusted R^2	0.457	0.454	0.453
F	11.7	10.9	11.1

Model (1) uses the log-transformed number of forward citations as the dependent variable. Model (2) uses the log-transformed number of applicant-excluded forward citations as the dependent variable. Finally, Model (3) uses the log-transformed number of self-citation-excluded forward citations as the dependent variable. All the models include patent and age fixed effects. Clustered Standard errors at the patent level are in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

6 Conclusions

The proliferation of patent assertion entities (PAEs) has become a topic of intense academic and policy debate. On the one hand, critics suggest that the PAE enforcement model imposes costs that are not proportionate to the value of the patented technology, while their litigation targets – often operating companies – have fewer defensive options since PAEs neither produce goods nor perform R&D: as a result, PAEs are responsible for a deadweight loss to the economy by discouraging operating companies from innovating. On the other hand, advocates of the PAE business stress that their patents are often stronger than those held by operating companies and that they serve as intermediaries in the market for invention.

The goal of this paper has been to enrich the debate by providing new evidence based on the patenting activity of PAEs in Europe, a region where the patent assertion landscape is growing rapidly and the imminent introduction of the Unified Patent Court (UPC) and the Unitary Patent (UP) are likely to be "game-changing events that could increase the amount of patent assertion activity in Europe" (Thumm, 2018).

By exploiting a unique database of patent transfers involving PAEs at the European Patent Office, we find that the presence of PAEs in Europe is not marginal. When considering only EP high-tech patents applied between 1997 and 2012, the share of those involving at least one PAE as either first applicant or buyer constitutes around 1.6% of the total. When focusing on patent transfers, the share of PAEs increases to around 2.9%.

Furthermore, we investigate the impact of PAEs' business model on innovation by looking at the pattern of citations received by patents acquired by PAEs. Building on the idea that patent citations are an indicator of the use of the protected technology by innovating and producing companies, we assume that a patent that stops being cited indicates that the protected technology is likely to no longer be used in further inventions. We thus firstly compare citation profiles of transferred patents, before and after the transfer, with citation profiles of never transferred patents, separating the former kind of patents in two groups, i.e. transferred to PAEs and transferred to operating companies (PEs). In a second step, we restrict our sample directly to transferred patents and we investigate whether there is a significant difference in terms of citation profiles between the ones transferred to PAEs and the ones transferred to PEs.

Our econometric results show that (1) PAEs acquire patents that are, on average, of high technological quality (compared to both never-transferred patents and patents transferred to PEs) and that are already in the declining phase of their technological life cycle; (2) after a transfer occurs, patents that are transferred to both PAEs and PEs do receive fewer citations; (3) the reduction in the number of citations received after the transfer is analogue for the two groups (PAEs and PEs), when we restrict the sample to transferred patents sharing very similar fixed, pre-transfer characteristics.

Is the typical PAE business model harmful for innovation processes? Looking at the effect of patent transfers on patent citation profiles, we conclude that PAEs do not behave as patent intermediaries: on average, citations decline faster for patents transferred to PAEs than for never transferred patents in the post-transfer period. However, PAEs may theoretically perform the socially valuable function of creating a "capital market for invention" by providing incentives for individual and small inventors, and making the patent market more liquid (McDonough, 2006; Myhrvold, 2010): by acquiring high-quality patents they may reward effective R&D efforts. In principle, the question is still open, although the fact that PAEs seem to transfer only a small fraction of their revenues to original patent inventors (Bessen et al., 2011) speaks in favor of an affirmative response.

In addition, our results also raise up a broader issue for the entire functioning of the market for technology. Indeed, at least for the ICT domain and for a specific kind of technologies (i.e. the ones targeted by PAEs), it seems that patent transfers occur mainly for strategic reasons, independently on the type of entity that purchases patents.

Our analysis is not without limitations. First, it would be worth bringing our study of PAEs and the patent intermediary activity closer to reality by adding information on licensing agreements to our setting. Indeed, we do not observe any transaction which does not involve patent sales. We thus unavoidably underestimate the PAE business in the patent market.

Second, we observe only patent transfers that occur during the granting process, again underestimating the presence of PAEs in the patent market. Observing data on patent transfers occurring after the grant would allow for a better understanding of the strategies pursued by PAEs to enter the market. PAEs are indeed often accused of buying and litigating patents as late as possible, when the unsuspecting infringers have already started the production of goods based on technologies protected by the patents concerned, so as to maximize licensing fees.

One last remark concerns the policy implications of our work. In order to keep PAEs from reducing innovation and to protect legitimate patent holders, some economists and legal scholars have recommended reforming national patent offices by requiring them to conduct an open review whenever a patent is sold or renewed (Barker, 2005) and, in general, to improve the quality of patents issued (Bradford and Durkin, 2012). While the former recommendation would be likely to increase the transparency of patent transactions, thereby reducing the incentives of opportunistic behaviors, the latter would instead probably be neutral with respect to PAEs' strategies. Indeed, while it is true that these policy reforms would reduce the number of weak patents issued – guaranteeing a more efficient market for intellectual property rights – it is also true that patents acquired by PAEs are on average not so weak. In all, by intervening on the entry-side of the market, there is the risk of reducing the incentives for all kind of intermediaries, with no clear consequences on the net efficiency of the whole IPR system.

References

Bacchiocchi, E. and Montobbio, F. (2010), 'International knowledge diffusion and home-bias effect: Do uspto and epo patent citations tell the same story?', *The Scandinavian Journal of Economics* **112**(3), 441–470.

Barker, D. G. (2005), 'Troll or no troll? policing patent usage with an open post-grant review', *Duke Law & Technology Review* **4**(1), 1–17.

Bertrand, M., Duflo, E. and Mullainathan, S. (2004), 'How much should we trust differences-in-differences estimates?', *The Quarterly Journal of Economics* **119**(1), 249–275.

Bessen, J. (2003), 'Patent thickets: Strategic patenting of complex technologies', *Research on Innovation Working Paper* .

Bessen, J., Ford, J. and Meurer, M. J. (2011), 'The private and social costs of patent trolls', *Regulation* **34**, 26.

Bessen, J. and Meurer, M. J. (2013), 'The direct costs from npe disputes', *Cornell L. Rev.* **99**, 387.

Blasnik, M. (2007), 'RECLINK: Stata module to probabilistically match records', Statistical Software Components, Boston College Department of Economics.
URL: *https://ideas.repec.org/c/boc/bocode/s456876.html*

Blind, K., Cremers, K. and Mueller, E. (2009), 'The influence of strategic patenting on companies' patent portfolios', *Research Policy* **38**(2), 428–436.

Bottazzi, L. and Peri, G. (2003), 'Innovation and spillovers in regions: Evidence from european patent data', *European economic review* **47**(4), 687–710.

Bradford, B. J. and Durkin, S. J. (2012), 'A proposal for mandatory patent reexaminations', *IDEA* **52**, 135.

Breschi, S., Lissoni, F. and Montobbio, F. (2005), 'The geography of knowledge spillovers: conceptual issues and measurement problems', *Clusters, networks and innovation* pp. 343–378.

Burk, D. L. and Lemley, M. A. (2009), *The patent crisis and how the courts can solve it*, University of Chicago Press.

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

Carter, J. B. (2013), 'Responding to a patent troll's threats', *ARK. LAW* p. 30.

Chien, C. V. (2008), 'Of trolls, davids, goliaths, and kings: Narratives and evidence in the litigation of high-tech patents', *NCL Rev.* **87**, 1571.

Coffano, M. and Tarasconi, G. (2014), 'Crios-patstat database: sources, contents and access rules', *Center for Research on Innovation, Organization and Strategy, CRIOS working paper* (1).

Cohen, L., Gurun, U. G. and Kominers, S. D. (2016), 'The growing problem of patent trolling', *Science* **352**(6285), 521–522.

Cohen, L., Gurun, U. and Kominers, S. D. (2014), 'Patent trolls: Evidence from targeted firms', *NBER Working Paper 20322* .

Cotropia, C. A., Kesan, J. P. and Schwartz, D. L. (2014), 'Unpacking patent assertion entities (paes)', *Minn. L. Rev.* **99**, 649.

Darts-IP (2018), Npe litigation in the european union, Technical report.

De Rassenfosse, G., Kracker, M. and Tarasconi, G. (2017), 'Getting started with patstat register', *Australian Economic Review* **50**(1), 110–120.

Elsten, C. M. and Hill, N. (2017), 'Intangible asset market value study?', *Journal of the Licensing Executives Society* **Volume LII No. 4**.

Feldman, R. (2012), *Rethinking patent law*, Cambridge, MA: Harvard University Press.

Feldman, R. and Frondorf, E. (2015), 'Patent demands and initial public offerings', *Stan. Tech. L. Rev.* **19**, 52.

Feng, J. and Jaravel, X. (2016), 'Who feeds the trolls? patent trolls and the patent examination process', *Working Paper* .

Fischer, T. and Henkel, J. (2012), 'Patent trolls on markets for technology. an empirical analysis of npes patent acquisitions', *Research Policy* **41**(9), 1519–1533.

Fischer, T. and Leidinger, J. (2014), 'Testing patent value indicators on directly observed patent value–an empirical analysis of ocean tomo patent auctions', *Research Policy* **43**(3), 519–529.

Fusco, S. (2013), 'Markets and patent enforcement: A comparative investigation of non-practicing entities in the united states and europe', *Mich. Telecomm. & Tech. L. Rev.* **20**, 439.

Fusco, S. (2016), 'Trips non-discrimination principle: Are alice and bilski really the end of npes', *Tex. Intell. Prop. LJ* **24**, 131.

Gans, J. S. and Stern, S. (2010), 'Is there a market for ideas?', *Industrial and Corporate Change* **19**(3), 805–837.

Griliches, Z. (1998), Patent statistics as economic indicators: A survey, Technical report.

Haber, S. H. and Werfel, S. H. (2015), 'Why do inventors sell to patent trolls? experimental evidence for the asymmetry hypothesis', *Stanford University Working Paper* .

Hagiu, A. and Yoffie, D. B. (2013), 'The new patent intermediaries: platforms, defensive aggregators, and super-aggregators', *The Journal of Economic Perspectives* **27**(1), 45–65.

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

Hall, B. H. and Ziedonis, R. H. (2001), 'The patent paradox revisited: An empirical study of patenting in the us semiconductor industry, 1979-1995', *Rand Journal of Economics* pp. 101–128.

Hoppe, H. C. and Ozdenoren, E. (2005), 'Intermediation in innovation', *International Journal of Industrial Organization* **23**(5), 483–503.

Jaffe, A. B. and Trajtenberg, M. (1999), 'International knowledge flows: evidence from patent citations', *Economics of Innovation and New Technology* **8**(1-2), 105–136.

Jaffe, A. B., 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.

Khan, B. Z. (2013), 'Trolls and other patent inventions: economic history and the patent controversy in the twenty-first century', *Geo. Mason L. Rev.* **21**, 825.

Kingston, W. (2001), 'Innovation needs patents reform', *Research Policy* **30**(3), 403–423.

Lanjouw, J. O. and Schankerman, M. (2004), 'Protecting intellectual property rights: Are small firms handicapped?', *The Journal of Law and Economics* **47**(1), 45–74.

Lechner, M. (2001*a*), 'Identification and estimation of causal effects of multiple treatments under the conditional independence assumption', *Econometric Evaluation of Labour Market Policies* pp. 43–58.

Leiponena, A. and Delcamp, H. (2018), 'The anatomy of a troll? patent licensing business models in the light of patent reassignment data', *Research Policy* **Forthcoming**.

Lemley, M. A. (2006), 'Should patent infringement require proof of copying', *Mich. L. Rev.* **105**, 1525.

Lemley, M. A. and Feldman, R. (2016), 'Patent licensing, technology transfer, and innovation', *The American Economic Review* **106**(5), 188–192.

Lemley, M. A. and Shapiro, C. (2006), 'Patent holdup and royalty stacking', *Tex. L. Rev.* **85**, 1991.

Leslie, C. R. (2008), 'Patents of damocles', *Ind. LJ* **83**, 133.

Lizzeri, A. (1999), 'Information revelation and certification intermediaries', *The RAND Journal of Economics* pp. 214–231.

Love, B. J. (2013), 'An Empirical Study of Patent Litigation Timing: Could a Patent Term Reduction Decimate Trolls Without Harming Innovators?', *University of Pennsylvania Law Review* pp. 1309–1359.

Love, B. J., Helmers, C., Gäßler, F. and Ernicke, M. (2017), *Patent Assertion Entities and Competition Policy*, D. Daniel Sokol edn, Cambridge University Press, chapter Patent Assertion Entities in Europe, pp. 104–129.

Lu, J. J. (2012), 'The Myths and Facts of Patent Troll and Excessive Payment: Have Nonpracticing Entities (NPEs) Been Overcompensated?', *Business Economics* **47**, 234–249.

Martínez, C. (2011), 'Patent families: When do different definitions really matter?', *Scientometrics* **86**(1), 39–63.

Maurseth, P. B. and Verspagen, B. (2002), 'Knowledge spillovers in europe: A patent citations analysis', *The Scandinavian Journal of Economics* **104**(4), 531–545.

Mayergoyz, A. (2009), 'Lessons from europe on how to tame us patent trolls', *Cornell Int'l LJ* **42**, 241.

McDonough, J. F. I. (2006), 'The myth of the patent troll: An alternative view of the function of patent dealers in an idea economy', *Emory LJ* **56**, 189.

Montobbio, F. and Sterzi, V. (2011), 'Inventing together: exploring the nature of international knowledge spillovers in latin america', *Journal of Evolutionary Economics* **1**(21), 53–89.

Myhrvold, N. (2010), 'The big idea: Funding eureka!', *Harvard Business Review* **88**(2), 40–50.

Noel, M. and Schankerman, M. (2013), 'Strategic patenting and software innovation', *The Journal of Industrial Economics* **61**(3), 481–520.

Orsenigo, L. and Sterzi, V. (2010), 'Comparative study of the use of patents in different industries', *KITES working paper 33/2010, Bocconi University* .

Ortiz, J. (2016), 'Non-practicing entities changing business model', *ClearIP View* **http://www.clearviewip.com/new-frontiers-npes/**.

Papst, D. (2012), 'Npes and patent aggregators. new, complementary business models for modern ip markets', *Licensing Journal, 32 (10)* pp. 1–5.

Parchomovsky, G. and Wagner, R. P. (2005), 'Patent portfolios', *University of Pennsylvania Law Review* **154**(1), 1–77.

Penin, J. (2012), 'Strategic uses of patents in markets for technology: A story of fabless firms, brokers and trolls', *Journal of Economic Behavior & Organization* **84**(2), 633–641.

Quinn, G. (2010), 'Mother of all patent trolls, acacia research, gets more funding', *IP WATCHDOG* .

Raffo, J. (2015), 'MATCHIT: Stata module to match two datasets based on similar text patterns', Statistical Software Components, Boston College Department of Economics.
URL: *https://ideas.repec.org/c/boc/bocode/s457992.html*

Reitzig, M., Henkel, J. and Heath, C. (2007), 'On sharks, trolls, and their patent prey–Unrealistic damage awards and firms' strategies of "being infringed"', *Research Policy* **36**(1), 134–154.

Risch, M. (2012), 'Patent troll myths', *Seton Hall L. Rev.* **42**, 457.

Rosenbaum, P. R. and Rubin, D. B. (1985), 'Constructing a control group using multivariate matched sampling methods that incorporate the propensity score', *The American Statistician* **39**(1), 33–38.

Schmoch, U. (2008), 'Concept of a technology classification for country comparisons', *Final report to the world intellectual property organisation (WIPO)* .

Schwartz, D. L. and Kesan, J. P. (2013), 'Analyzing the role of non-practicing entities in the patent system', *Cornell L. Rev.* **99**, 425.

Scott Morton, F. and Shapiro, C. (2016), 'Patent assertions: Are we any closer to aligning reward to contribution?', *Innovation Policy and the Economy* **16**(1), 89–133.

Serrano, C. J. (2010), 'The dynamics of the transfer and renewal of patents', *The RAND Journal of Economics* **41**(4), 686–708.

Serrano, C. J. and Ziedonis, R. (2018), How redeployable are patent assets? evidence from failed startups, Technical report, National Bureau of Economic Research.

Shapiro, C. and Scott-Morton, F. (2014), 'Strategic patent acquisitions', *Antitrust Law Journal* **2**, 463–499.

Shrestha, S. K. (2010), 'Trolls or market makers? an empirical analysis of nonpracticing entities', *Columbia Law Review* **110**, 114–160.

Squicciarini, M., Dernis, H. and Criscuolo, C. (2013), 'Measuring patent quality: Indicators of technological and economic value', *OECD Science, Technology and Industry Working Papers* .

Thumm, N. (2018), 'The good, the bad and the ugly–the future of patent assertion entities in europe', *Technology Analysis & Strategic Management* pp. 1–11.

Torrisi, S., Gambardella, A., Giuri, P., Harhoff, D., Hoisl, K. and Mariani, M. (2016), 'Used, blocking and sleeping patents: Empirical evidence from a large-scale inventor survey', *Research policy* **45**(7), 1374–1385.

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

Trajtenberg, M., Henderson, R. and Jaffe, A. (1997), 'University Versus Corporate Patents: A Window On The Basicness Of Invention', *Economics of Innovation and New Technology* **5**(1), 19–50.

Tucker, C. E. (2014), 'Patent trolls and technology diffusion: The case of medical imaging', *Working Paper SSRN 1976593* .

Turner, J. L. (2011), 'Input complementarity, patent trolls and unproductive entrepreneurship', *Working Paper SSRN 1916798* .

Walsh, J. P., Lee, Y.-N. and Jung, T. (2016), 'Win, lose or draw? the fate of patented inventions', *Research Policy* **45**(7), 1362–1373.

Wasi, N. and Flaaen, A. (2015), 'Record linkage using Stata: Preprocessing, linking, and reviewing utilities', *Stata Journal* **15**(3), 672–697.

Yeh, B. T. (2013), 'An overview of the" patent trolls" debate', *CRS Report for Congress* .

Appendixes

A1. Harmonization and disambiguation of applicant names listed in the EP Register

EPR assigns to each recorded applicant a unique internal identifier based on a combination of the fields "NAME" and "ADDRESS". As stressed above, several incongruences could emerge in identifying patent-ownership changes, essentially because applicants' identities have not been harmonized or disambiguated before being listed in the EPR database (De Rassenfosse et al., 2017). Indeed, if the same applicant changes name and/or address in its patenting life, the database will automatically assign a new event for all the patents it owns, with a new identifier attached (without updating the former one). Similarly, if the same applicant owns two patents, but the name and/or the address have been recorded differently in the two original documents (*i.e.* due to typing errors or different abbreviations), two different identifiers will be assigned accordingly. These incongruences thus represent a relevant source of bias when analyzing changes in patent ownership and when matching this source of data with external information.

To overcome this issue and partially reduce the number of false positives when analyzing patent legal events, we harmonize and standardize applicants' names and addresses. Since original EP-Register data on applicants' names and addresses come in a text string, we pre-process the data as follows:

1. *Parsing, cleaning and standardizing applicants' names.* The original text string for applicants' names is parsed into relevant sub-components, cleaned by removing special characters and stop words, and standardized with respect to abbreviations for business entities. In this step, we apply the STATA utility "stnd_compname" (Wasi and Flaaen, 2015).[50]

2. *Parsing, cleaning and standardizing applicants' addresses.* Similarly, the original text string for applicants' addresses is parsed into rele-

[50] We extend the standardization procedure proposed by Wasi and Flaaen (2015) by extending the list of abbreviations for company names to countries different from the US. Precisely, we add abbreviations usually appearing in Germany, the UK, France, Italy, Spain, Denmark, Sweden, Switzerland, Finland, Russia, and Japan.

vant sub-components, cleaned by removing special characters and stop words, and standardized for abbreviations. In this step, we apply the STATA utility "stnd_address" (Wasi and Flaaen, 2015). Moreover, we isolate and standardize the country field from the applicant's address.

Based on cleaned names and addresses, we re-assign to each original applicant four new internal identifiers, following different rules: *i)* a new identifier grouping applicants showing the same name and complete address (*id_name_and_address*); *ii)* a new identifier grouping applicants sharing the same name and the same country (*id_name_and_country*); *iii)* a new identifier grouping applicants showing the same name, independently from the address (*id_name*); *iv)* a new identifier grouping applicants showing the same complete address, independently from the name (*id_address*). The first identifier follows the same logic adopted by EPR, but it groups original applicants more precisely than EPR does (reducing the number of false positives and still minimizing the number of false negatives). Conversely, the other identifiers go for higher recall, but at the cost of being less precise (they reduce the number of false negatives, but at the cost of allowing for higher numbers of false positives).

Further refinement: String similarity within parties involved in possible patent transfers

As a second step performed to augment the precision in capturing patent transfers, we directly focus on patents showing potential transaction events during the granting phase. Here we look firstly at the applicant addresses reported in the patent document (*id_address* above). If two applicants listed in the same document share the same address, we consider them as one entity. Then, within the rest of the applicant names listed in the same patent document, we apply the STATA tool MATCHIT (Raffo, 2015) to assign a probability that two unique parties are actually the same. More precisely, MATCHIT is a tool developed to join observations from two datasets based on string variables which do not necessarily need to be exactly the same. It allows for a fuzzy similarity between two different text variables. We consider two unique entities as the same according to three different similarity thresholds that are selected so as to minimize possible errors (0.9, 0.95 and 0.99). Coherently, we assign three new internal IDs to all the applicants listed in the EPR database.

According to the methodology described above, we end up with a final sample of 460,895 unique applicants (fixing the similarity threshold at 0.95). This means a reduction in the raw number of unique customer IDs of around 36%. Looking at the transfers individuated, they are responsible for 369,828 patents with at least one change in the applicant field recorded (12.8% of the total number of applications registered at the EPO). According to the raw information on the applicant identifiers listed in the original data, patent possibly transferred were 700,954. Therefore, for 331,126 applications, the potential transfer emerging from the raw data is simply a change in the applicant name or address. Table 6 reports the number of patent transfers according to the three similarity thresholds applied.[51]

Table 11: Applicant names consolidation in EP Register: Number of transfers

Matchit Threshold	**EPR Raw**	**EPR Cleaned**	**Reduction**
0.9	700,954	364,032	-48.07%
0.95	700,954	369,828	-47.24%
0.99	700,954	371,107	-47.06%

The Table reports the reduction in the number of potential patent transfers individuated at EPO once consolidated the applicant identities.

[51]It is worth to notice that we discard from the analysis international applications that have not entered the EP regional phase. Indeed, no EP publication exists in the PATSTAT core tables for these applications, and we cannot retrieve information about those patents.

A2. Matching quality

A2.1. Matching quality (1): transferred patents vs never transferred patents

This Section will describe the quality of the matching methods we implemented to perform the conditional DD estimates presented in Section 4.3. First, we check whether the common support condition holds. This condition ensures that we estimate only effects in regions where two observations, one belonging to the treated and the other to the control group, can have a similar participation probability. Figure 4 displays a graphic analysis of the Kernel density distribution for the two groups, before the implementation of the matching.[52] Though the shape of the two distributions differs, there is a large overlap between the distribution of the propensity score of the treated and the control group, ensuring that the common support condition holds.

Second, we check whether the matching on the propensity score actually manages to balance the distribution of the relevant variables in the control and the treatment groups. The literature suggests several methods to evaluate the matching quality. A common methodology, first introduced by Rosenbaum and Rubin (1985), is the two-sample t-test to check for significant differences in covariate means, for both groups, before and after the matching. Table 12 reports the t-test for all the covariates we included in the probit regression to estimate the propensity score for the unmatched and the matched samples.

Before the matching, there is a significant difference in the mean between the treated and the control group for all the variables we are interested in (with the exception of the number of backward citations). However, all these differences are no longer statistically significant after implementing the matching procedure, confirming its good performance in balancing the covariates.

Furthermore, to asses the size of the bias reduction obtained through the propensity score matching method we compute the standardized bias and we compare its size before and after the matching (Rosenbaum and Rubin,

[52] Lechner (2001*a*) argues that it is possible to assess the overlap between sub-samples through a graphic analysis of the propensity score density distribution for the treated and the control group, before the matching.

Figure 4: Kernel density distributions of the propensity score before the matching

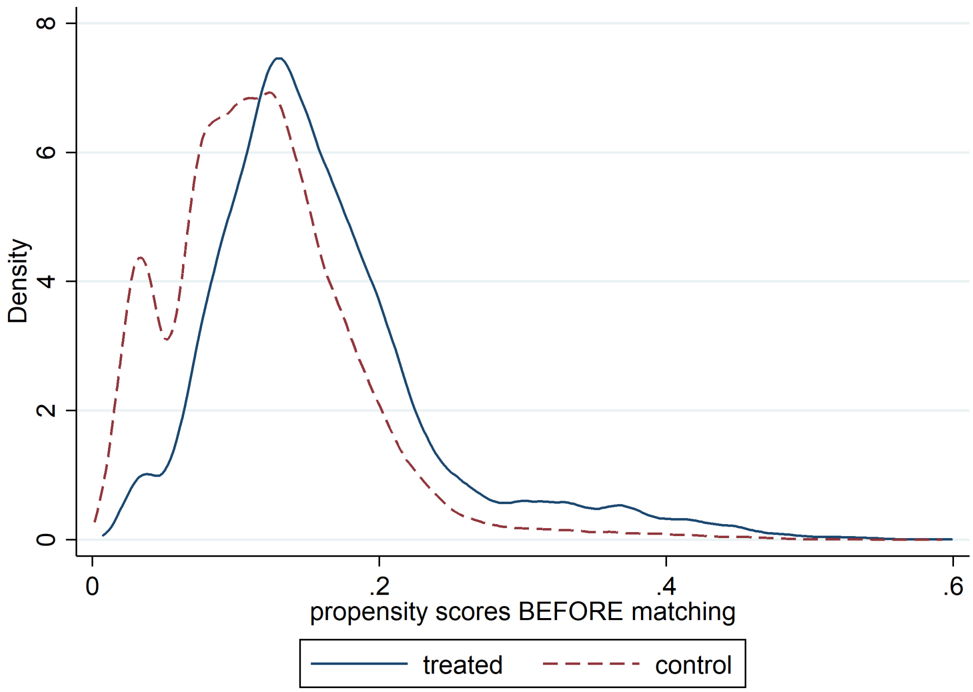

Table 12: Descriptive statistics for the Unmatched and the Matched Sample

Variable	**Unmatched (U)** **Matched (M)**	**Mean**		**%reduct**		**t-test**	
		Treated	*Control*	*%bias*	*|bias|*	*t*	*p>t*
AVG 4 YEAR CITS	U	.44294	.39819	9.2		12.41	0.000
	M	.44285	.44195	0.2	98.0	0.18	0.860
CLAIMS (LN)	U	25.329	25.134	3.3		4.38	0.000
	M	25.328	2.535	-0.4	88.7	-0.36	0.722
TEAM SIZE (LN)	U	.84266	.79486	7.6		9.97	0.000
	M	.84254	.84022	0.4	95.2	0.36	0.718
ORIGINALITY	U	.70487	.68984	7.0		9.00	0.000
	M	.70486	.70482	0.0	99.7	0.02	0.984
BACKWARD CITS	U	1.6473	1.6452	0.4		0.50	0.617
	M	1.6473	1.6421	1.0	-149.2	0.93	0.354
PATENT STOCK (LN)	U	.65593	.99809	-15.6		-18.94	0.000
	M	.656	.62192	1.6	90.0	1.73	0.083
INDIVIDUAL	U	.01973	.01332	5.0		7.07	0.000
	M	.01973	.02	-0.2	95.8	-0.19	0.850
COAPPLICANT	U	.11117	.03722	28.5		46.05	0.000
	M	.11108	.10854	1.0	96.6	0.80	0.425

Dummies for patent age, year of filing, inventor's country of residence and technological fields included in the probit model.

1985). Table 6 reports the mean and the median standardized bias, before and after the matching. Though there is no clear threshold under which it is possible to tell the success of the matching procedure with certainty, a bias reduction below 3 or 5 per cent is generally considered as sufficient (Caliendo and Kopeinig, 2008). As the Table shows, both the mean and the median standardized biases fall below the one per cent level after the matching, confirming the reliability of the matching on the propensity score.

Table 13: Mean and median standardized bias for the matched and unmatched sample

Sample	Mean Bias	Median Bias
Unmatched	9.6	7.3
Matched	0.6	0.4

Finally, since intuitively the matching procedure is implemented to "correct" for differences in terms of the probability of receiving the treatment between the treated and the control group, we can look at the visual representation of the propensity score distributions and make a comparison before and after the matching. As Figure 5 displays, the difference in the Kernel density distribution of the estimated propensity scores abundantly diminishes with respect to the pre-matching situation offered by Figure 4: the two distributions almost perfectly overlap, once again suggesting that the propensity score matching procedure successfully corrects for the selection on observable factors.

We present the results from the probit regression implemented for calculating the propensity scores in Table 14. The probability of a patent being transferred is positively correlated with the average number of yearly citations it receives during the 4 years after the filing, with the team size, with the patent originality, with the number of citations made and with the dummy coapplicant. Conversely, a patent transfer is negatively correlated with the size of the applicant patent portfolio: small entities are more likely than large entities to sell patents. Moreover, it is negatively associated with the patent scope, proxied by the number of claims. The dummy individual does not show a significant effect.

Figure 5: Kernel density distributions of the propensity score after the matching

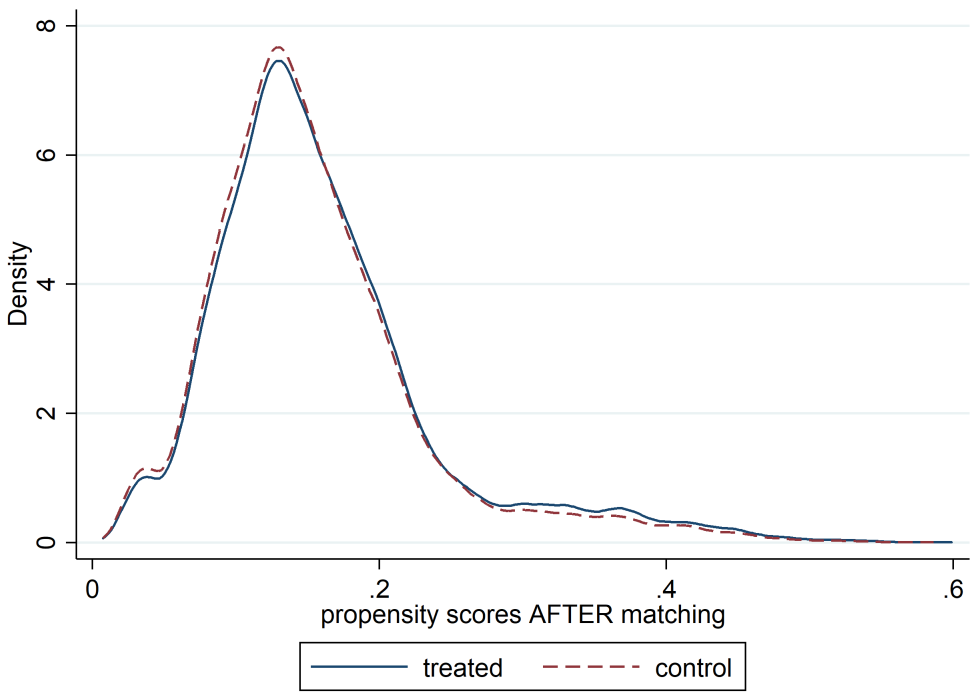

Table 14: Probit results

	(1)
AVG 4 YEAR CITS	0.0454***
	(0.0091)
CLAIMS (LN)	-0.0188**
	(0.0074)
TEAM SIZE (LN)	0.0137*
	(0.0070)
ORIGINALITY	0.0443**
	(0.0221)
BACKWARD CITS	0.0646***
	(0.0086)
PATENT STOCK (LN)	-0.0389***
	(0.0020)
INDIVIDUAL	0.0543
	(0.0336)
COAPPLICANT	0.6337***
	(0.0169)
FILING YEAR DUMMIES	yes
AGE DUMMIES	yes
TECHNOLOGY DUMMIES	yes
COUNTRY DUMMIES	yes
CONSTANT	-1.036***
OBSERVATIONS	155,943

Standard errors in parenthesis. Dependent Variable: probability of a patent transfer. For the description of the variables included and for their discussion see Section 4.3. * $p < .1$, ** $p < .05$, *** $p < .01$.

A2.2. Matching quality (2): transferred patents to PAEs vs transferred patents to PEs

In this Section we provide the analysis of the quality of the matching methods we implemented to perform the conditional DD estimates described in Section 4.4. To test for the quality of the matching we follow the same schema as before. It is worth noticing here that our two groups of interest are now composed by, respectively, patents transferred to PAEs (treated group) and patents transferred to PEs (control group). The treatment for our observed patents is thus represented by being transferred to a PAE.

First, we visually reproduce the Kernel density distribution for the two groups, before and after the implementation of the matching (Fig. 6 and 7). Though the shape of the two distributions differs before performing the matching, there is a sufficient overlap between the distribution of the propensity score of the treated and the control groups, ensuring that the common support condition holds. After the match, the difference in the Kernel density distribution of the estimated propensity scores abundantly diminishes with respect to the pre-matching situation and the two distributions almost perfectly overlap. This first visual test goes in the direction of confirming that our matching procedure performs properly.

Second, we run the two-sample t-test to check for significant differences in covariate means, for both groups, before and after the matching. Table 15 reports the t-test for all the covariates we included in the probit regression to estimate the propensity score for the unmatched and the matched samples. Since in this last part of the empirical analysis we focus only on transferred patents, it is worth noticing that now we also include the year of the transfer within our set of covariates.

Before the matching, there is a significant difference in the mean between the treated and the control group for all the variables we are interested in. However, all these differences are no longer statistically significant after implementing the matching procedure, confirming its good performance in balancing the covariates.

A final test is performed to asses the size of the bias reduction obtained through the propensity score matching method. To do so, we compute the standardized bias and we compare its size before and after the matching. Results are reported in Table 16. The mean and the median standardized

Figure 6: Kernel density distributions of the propensity score before and after the matching (2)

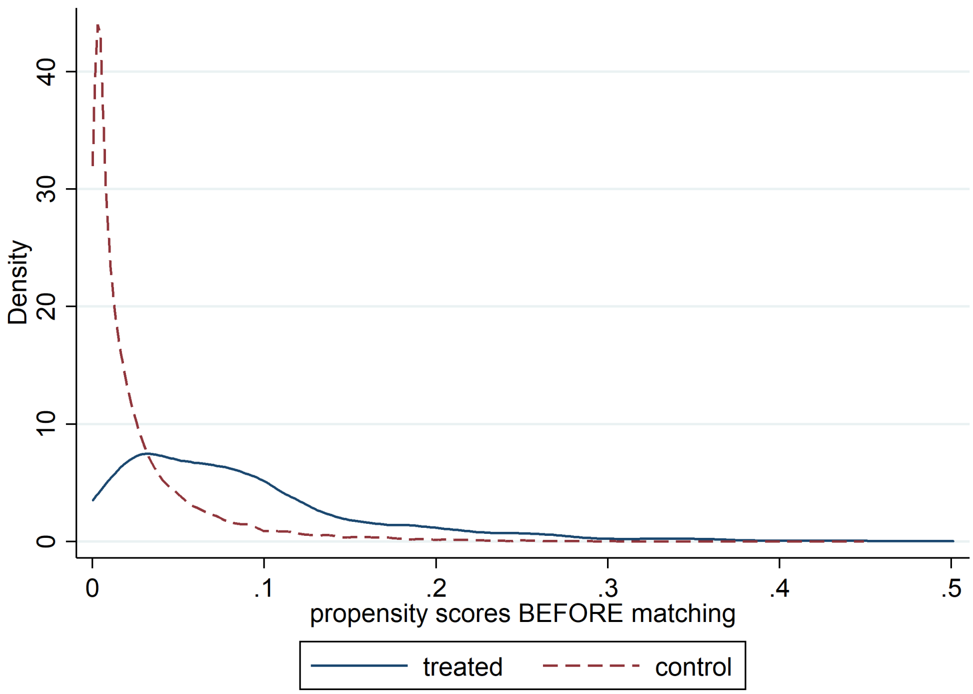

Figure 7: Kernel density distributions of the propensity score before and after the matching (2)

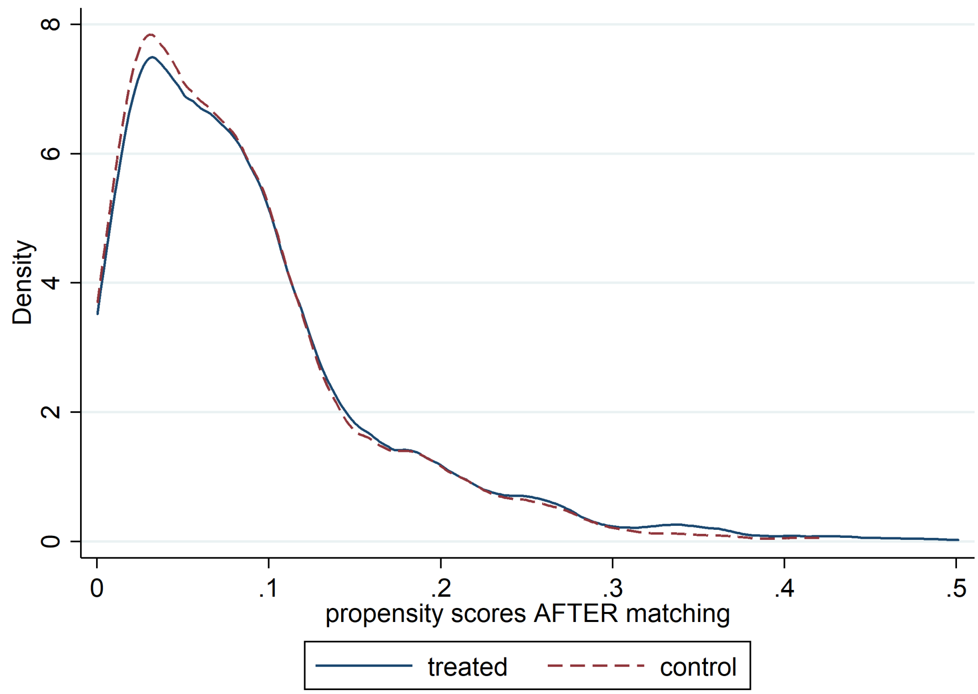

Table 15: Descriptive statistics for Unmatched and Matched Sample (2)

Variable	**Unmatched (U) Matched (M)**	**Mean**		**%reduct**		**t-test**	
		Treated	*Control*	*%bias*	*|bias|*	*t*	*p>t*
AVG 4 YEAR CITS	U	.58816	.41841	31.3		10.15	0.000
	M	.58651	.57204	2.7	91.5	0.51	0.611
CLAIMS (LN)	U	2.6265	2.5272	16.5		4.89	0.000
	M	2.6249	2.6296	-0.8	95.3	-0.16	0.872
TEAM SIZE (LN)	U	.76791	.83147	-10.2		-2.93	0.003
	M	.77057	.78367	-2.1	79.4	-0.45	0.655
ORIGINALITY	U	.7311	.69621	17.9		4.80	0.000
	M	.73064	.73212	-0.8	95.8	-0.17	0.865
BACKWARD CITS	U	1.6697	1.627	8.0		2.35	0.019
	M	1.6683	1.6724	-0.8	90.3	-0.16	0.874
PATENT STOCK (LN)	U	.47993	.8142	-17.0		-4.55	0.000
	M	.48159	.48033	0.1	99.6	0.02	0.988
INDIVIDUAL	U	.0069	.02161	-12.4		-2.97	0.003
	M	.00693	.003	3.3	73.3	1.16	0.245
COAPPLICANT	U	.01956	.11319	-38.3		-8.69	0.000
	M	.01963	.01894	0.3	99.3	0.10	0.917
POOL	U	.44304	.55081	-21.7		-6.29	0.000
	M	.44457	.43464	2.0	90.8	0.42	0.677

Dummies for year of filing, year of transfer, inventor's country of residence and technological fields included in the probit model.

biases fall below, respectively, the two and the one per cent level after the matching, confirming the reliability of the matching on the propensity score.

Table 16: Mean and median standardized bias for the matched and unmatched sample (2)

Sample	Mean Bias	Median Bias
Unmatched	19.2	17.0
Matched	1.4	0.8

Finally, we report the results from the probit regression implemented for calculating the propensity scores in Table 17. The probability of a patent being transferred to a PAE is positively correlated with the average number of yearly citations it receives during the 4 years from the filing and with the number of claims. Conversely, a patent transfer to a PAE is negatively correlated with the inventor team size, the dummy pool, the dummy coapplicant and the dummy individual.

Table 17: Probit results

	(1)
AVG 4 YEAR CITS	0.1646***
	(0.0311)
CLAIMS (LN)	0.0682**
	(0.0289)
TEAM SIZE (LN)	-0.0723***
	(0.0277)
ORIGINALITY	0.2014**
	(0.0968)
BACKWARD CITS	0.0162
	(0.0349)
PATENT STOCK (LN)	-0.0113
	(0.0094)
INDIVIDUAL	-0.5233***
	(0.1770)
COAPPLICANT	-0.7290***
	(0.0952)
POOL	-0.2823***
	(0.0358)
FILING YEAR DUMMIES	yes
TRANSFER YEAR DUMMIES	yes
TECHNOLOGY DUMMIES	yes
COUNTRY DUMMIES	yes
CONSTANT	-5.4344
OBSERVATIONS	29,890

Standard errors in parenthesis. Dependent Variable: probability of a patent transfer to a PAE. For the description of the variables included and for their discussion see Section 4.3 and Section 4.4. * $p < .1$, ** $p < .05$, *** $p < .01$.

A3. Robustness Checks

Table 18: Baseline Models: Negative Binomial results

	(1) RAW	(2) RAW	(4) NO-APP	(5) NO-APP
PE	0.056***	0.052***	0.053***	0.051***
	(0.013)	(0.013)	(0.013)	(0.013)
PAE	0.35***	0.34***	0.35***	0.34***
	(0.053)	(0.051)	(0.053)	(0.051)
TRADED	0.041**	0.039**	0.038**	0.036**
	(0.017)	(0.016)	(0.016)	(0.016)
TRADED*PAE	-0.34***	-0.30***	-0.31***	-0.27***
	(0.065)	(0.066)	(0.065)	(0.065)
TEAM SIZE (LN)		0.33***		0.33***
		(0.011)		(0.011)
ORIGINALITY		0.43***		0.41***
		(0.021)		(0.021)
CLAIMS (LN)		0.37***		0.37***
		(0.0084)		(0.0081)
COAPPLICANT		0.0090		0.0025
		(0.026)		(0.024)
INDIVIDUAL		-0.016		-0.021
		(0.049)		(0.045)
PATENT STOCK (LN)		0.0014		0.00090
		(0.0015)		(0.0015)
Age FE	Yes	Yes	Yes	Yes
Filing Year FE	Yes	Yes	Yes	Yes
Technology FE	Yes	Yes	Yes	Yes
Country FE	Yes	Yes	Yes	Yes
Patent FE	No	No	No	No
Observations	2154839	2083259	2154839	2083259
Pseudo R^2	0.047	0.055	0.049	0.057

Models 1-3 estimate the effect on the raw count of forward citations. Models 4-6 use the count of applicant-excluded citations as the dependent variable. Clustered Standard errors at the patent level are in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table 19: Baseline models (Exclusion of self citations)

	(1)	(2)	(3)	(4)
	Cit (LN)	**Cit (LN)**	**Cit (LN)**	**Cit (LN)**
PE	0.053***	0.051***	0.013***	0.014***
	(0.0030)	(0.0030)	(0.0030)	(0.0030)
PAE	0.13***	0.18***	0.11***	0.11***
	(0.013)	(0.018)	(0.017)	(0.018)
TRADED	-0.040***	-0.037***	-0.00064	-0.0010
	(0.0030)	(0.0030)	(0.0029)	(0.0030)
TRADED*PAE		-0.10***	-0.096***	-0.095***
		(0.017)	(0.016)	(0.016)
TEAM SIZE (LN)				0.050***
				(0.0019)
ORIGINALITY				0.082***
				(0.0032)
CLAIMS (LN)				0.078***
				(0.0014)
COAPPLICANT				-0.0020
				(0.0037)
INDIVIDUAL				-0.011*
				(0.0063)
PATENT STOCK (LN)				0.0023***
				(0.00027)
Age FE	Yes	Yes	Yes	Yes
Filing Year FE	No	No	Yes	Yes
Technology FE	No	No	Yes	Yes
Country FE	No	No	Yes	Yes
Patent FE	No	No	No	No
Observations	2,154,839	2,154,839	2,154,839	2,154,839
Number of patents	178,564	178,564	178,564	178,564
Adjusted R^2	0.047	0.047	0.087	0.099
F	120.2	94.3	103.0	198.2

All the models use the count of self-citation-excluded forward citations as the dependent variable. Column (1) reports our most parsimonious estimation without our interaction of interest, with only patent age fixed effects included. In column (2) we add our interaction of interest. In column (4) we also control for a series of dummies: patent filing year, inventor's country of residence and technological domain. Column (4) is our preferred specification in which we add the full set of covariates. In column (5) we include patent fixed effects (and exclude time invariant controls). Clustered Standard errors at the patent level are in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$