Research Policy 51 (2022) 104590

Contents lists available at ScienceDirect

Research Policy

journal homepage: www.elsevier.com/locate/respol

Overcoming inefficiencies in patent licensing: A method to assess patent essentiality for technical standards

Rudi Bekkers [a,*], Elena M. Tur [a], Joachim Henkel [b], Tommy van der Vorst [c], Menno Driesse [c], Jorge L. Contreras [d]

[a] *School of Innovation Sciences, Eindhoven University of Technology, Netherlands*
[b] *Technical University of Munich, Germany*
[c] *Dialogic innovation & interaction, Netherlands*
[d] *University of Utah, United States of America*

ARTICLE INFO

JEL classifications:
O34
O38
O32
L15
L24

Keywords:
Standard essential patents
Essentiality
Patent licensing
Technical standards

ABSTRACT

The market for patent licenses, despite its paramount importance for technological innovation, has various inefficiencies. A particular problem with widely used technical standards such as LTE and Wi-Fi is the lack of information regarding which patents are "essential" to implement the standard. This information is crucial because it simplifies determining infringement and implies specific "FRAND" licensing rules. While many standards-developing organizations stipulate that such patents are explicitly declared, little is known about which are actually essential. The absence of publicly available information on essentiality incurs significant social costs due to the resulting friction in the licensing market. With the growing use and importance of standards to mobility and energy markets, and to the Internet of Things, these costs are likely to rise. Responding to calls from industry, courts and policymakers, commercial and academic studies have attempted to assess essentiality, but they all have limitations. This paper reports on the technical feasibility of a system of expert assessments for patent essentiality. Based on a factorial design, we conducted a field experiment with 20 patent examiners performing over 100 assessments. Comparing the outcomes to a high-quality reference point shows that sufficiently accurate expert assessments, at a price level that allows large scale testing, are technically feasible, and we identify routes to further improvement.

1. Introduction

In the field of high-tech products, academic studies focus a great deal not only on the product market but also on the markets for knowledge and related patents (Arora et al., 2001a, 2001b). An extensive body of literature has emerged on the patents required to implement technical standards, which are critical for a broad range of industries from telecommunications and computing to transportation, energy and manufacturing, even more so due to the spread of the wireless Internet of Things. Such "standard-essential patents" (SEPs) are of particular interest: unlike other patents, it is not possible for a party making products that incorporate these standards to design around such patents, thus creating an unusually strong bargaining position for patent owners. Given such a scenario and the societal importance of technical standards, the question this Special Issue asks – "Are intellectual property rights working for society?" – is particularly relevant in the context of SEPs.

Implementors of such standards often face thousands of patents that their owners claim to be potentially standard essential. The market for licenses for these patents shows serious imperfections, related to (among other things) transaction costs and information asymmetry. There are more and more calls for increased transparency to address these market inefficiencies. The European Commission (EC), in its 2017 communication on Standard Essential Patents, stated: "*There is therefore a need for a higher degree of scrutiny on essentiality claims. This would require scrutiny being performed by an independent party with technical capabilities and market recognition, at the right point in time*" (European Commission, 2017). More recently, the EC went one step further and its Executive Vice-President and Commissioner for Competition, Ms. Margrethe Vestager, announced that "*the Commission will explore setting up an independent system of third-party essentiality checks in view of improving legal certainty and reducing litigation costs*" (European Commission, 2020b).

* Corresponding author.
E-mail address: r.n.a.bekkers@tue.nl (R. Bekkers).

https://doi.org/10.1016/j.respol.2022.104590
Received 27 July 2021; Received in revised form 24 June 2022; Accepted 26 June 2022
Available online 18 July 2022

Insights into how to design such a system are therefore of high societal interest.

The Commission also recognizes that such a system must be balanced against the costs. While highly accurate assessments for single patents are feasible if cost and time are not an issue (as is often the case for single patents examined in the context of litigation), a large-scale system that systematically assesses the essentiality of patents for a given standard, requiring highly qualified individuals, would necessitate correspondingly high costs and time. Without pre-empting who would finance – or co-finance – such a large-scale system, and these parties' willingness (and self-interest) to budget for this, we are talking about potentially thousands of patents in certain patent-intensive industries such as mobile and wireless communications, and therefore assume any financier will be, to some degree, sensitive to cost. Arguably even more important is capacity: these assessment tasks require highly qualified reviewers, hence resources are limited. Should thousands of patents be assessed in a reasonable amount of time, then time efficiency is paramount.

Motivated by the above, this paper investigates whether essentiality assessments can be sufficiently efficient (in terms of time and costs) as well as sufficiently accurate, to enable the establishment of a large-scale system of such assessments, thus overcoming important inefficiencies in the SEP licensing market. Considering the complexity of the assessment task, but also the need for efficiency, this paper reports on assessments carried out in approximately 8 h on average (i.e., a one-day budget), acknowledging that some existing assessments, while considered accurate, utilize five to ten times more resources (see Section 3.2 below).

We formulated the following research questions:

Q1. How accurately can qualified assessors determine the essentiality of a given patent for a specific standard, and how much time do they take for this task?

Q2. How does the availability of claim charts[1] and the chosen definition of essentiality affect the above relationship?

Note that we focus on the "technical feasibility" of a large-scale system of essentiality testing. For questions concerning "institutional feasibility", such as who should set up such a system, who should carry out the assessments, who would finance it, and whether there is sufficient support from stakeholders, we refer to the complementary work presented in Bekkers et al. (2020a). While we focus on patent essentiality, we acknowledge that patent validity, enforceability, infringement and value are important aspects in terms of transparency and licensing discussions, and sometimes relate to essentiality,[2] but are beyond the scope of this study.

In a field experiment based on a factorial design, twenty experts, all patent examiners, performed 109 assessments in just over 100 working days. We involved patent examiners not only because they are qualified for this task, but also because of their impartiality, not being driven by political or economic motives. Using patent pool data as a high-quality reference point, we obtained not only data on patents accepted by pools, but also on those rejected by pools and the full information set provided at that time to pools – including claim charts prepared by the patent owner. This scenario enabled us to replicate the pool's assessment procedure as accurately as possible. Based on our analysis, we conclude that sufficiently accurate assessments, at a price level that allows large scale testing, are technically feasible, and we identify routes to further improvement.

In addition to addressing an important issue for standard-essential patents, this study has implications for the patent system in general. The patent system is built on an examination and issuance process that inherently involves trade-offs between accuracy and cost. Devoting greater resources to patent examination *ex ante* may result in the issuance of fewer patents eventually found to be invalid in litigation (*ex post*), but hardly any are ultimately challenged in court. Patent offices make calculated decisions regarding the level of resources to commit to examining each patent, knowing that complete accuracy is not possible. In order to improve accuracy without incurring excessive costs, experiments have been conducted to enlist the assistance of external agents in patent examination. One such pilot program called "Peer to Patent" was conducted by the USPTO and New York Law School between 2007 and 2011. This program allowed "citizen-experts" to review selected patent applications (mostly relating to computing, software and business methods), to identify and rate prior art, and to offer other input to the examination process. The study described in this article thus contributes to the literature regarding *ex ante* improvements to patent-based examination systems, and by focusing on essentiality rather than validity, it complements the above efforts. Our study's insights may be useful for designing future experiments that explore these processes.

The main contribution in this paper is empirical and policy related. Complementing this, we also make a theoretical contribution. We show that the notice function of patents (e.g., Bessen and Meurer, 2009), commonly a private good for the patent owner and the maker of an infringing product, takes on a public-good character in the case of SEPs. This insight justifies public ordering.

In the following Section 2, we discuss market inefficiencies in the licensing of patents, focussing on SEPs. Section 3 looks at existing attempts to determine essentiality, including patent pools, commercial and academic studies, and court case analyses, and we review recent endeavours using Artificial Intelligence (AI) approaches. Section 4 presents the experiment design and data for our research, followed by the results in Section 5, then conclusions and a discussion in Section 6.

2. Market inefficiencies in licensing standard-essential patents

2.1. The imperfect market for patent licenses

As economists have agreed since Adam Smith (1774), well-functioning markets increase efficiency through specialization and division of labour. However, efficient markets require parties to have full information. This assumption is frequently violated by information asymmetry, a lack of transparency, and uncertainty about future events.

A challenge with trading knowledge is that in many cases it can easily be expropriated. Intellectual property rights, and patents in particular, can alleviate this problem by increasing the appropriability of knowledge (Lamoreaux and Sokoloff, 1999; Arora et al., 2001a, 2001b; Arora and Ceccagnoli, 2006; Gans et al., 2008). Accordingly, the market for patent licenses (including, for the sake of simplicity, both patent licenses and assignments) has acquired a high importance in its own right (Madiès et al., 2014).

The market for patent licenses shows various imperfections due to transaction costs, both in terms of motivation and coordination costs (Milgrom and Roberts, 1992: 29). Motivation costs cause inefficiencies through information asymmetry (Caves et al., 1983), which is typically bilateral in the market for patent licenses. The patent owner might have private information about the value of the invention, its limitations, or required complementary knowledge not evident from the patent text. In turn, the prospective licensee will generally know more about the

[1] In this paper, "claim chart" refers to a document that the patent owner produces to map, in detail, the claim(s) and claim feature(s) in the patent to those sections of the standard document for which they are believed to be essential.

[2] The relationship between essentiality and validity can be complex. A broad interpretation of a patent's claims may lead to a patent being essential; however, that same breadth might actually render the patent invalid. Patent essentiality and infringement are related but not identical. Many standards define the behaviour of different device categories (for instance mobile phones vs. network switching equipment, and CD discs vs. CD players). A patent that is essential to the standard may in fact only be required for network switching equipment and will not be infringed by a mobile phone that fully complies with the standard. Moreover, standards often define optional features. While many SDOs have definitions of essentiality that include patents required for optional features, an implementer who decides to exclude an optional feature will not infringe the essential patents relating to this feature.

invention's potential applications and resulting economic value. In addition, there is often uncertainty for both parties regarding the eventual grant and scope of pending applications, the legal stability of patents, or infringement of a patent by a given product. There is also a lack of transparency due to the sheer number of patents. To some extent, the consequences of such information asymmetries and uncertainties can be mitigated by a suitable choice of licensing terms (Gallini and Wright, 1990; Beggs, 1992; Macho-Stadler et al., 1996), but inefficiencies remain.

The inefficiencies in the market for patent licenses can have serious welfare consequences, including the under-utilization of existing technologies, inadvertent or intentional infringement, and unfounded or, in the case of inadvertent infringement and subsequent lock-in, excessive royalty demands (e.g., Jaffe and Lerner, 2006; Lemley and Shapiro, 2007). These consequences are particularly severe when it comes to widely used, standardized technologies, as we discuss next.

2.2. Specific problems posed by licensing standard essential patents

A significant amount of product categories in the ICT sector utilises interoperability standards. Every smartphone, tablet and laptop sold today incorporates hundreds of different standards (Biddle et al., 2010). Likewise, hundreds of different firms and research organizations collaborate to develop widely deployed telecommunications standards such as UMTS (3G), LTE (4G), 5G, and Wi-Fi. At the same time, these standards are implemented by a large number of equipment makers. The rise of the Internet of Things (IoT) and smart technologies is expected to broaden the application of these technologies to domains like transportation, energy and manufacturing, thereby significantly increasing the number of implementers.

Interoperability standards are typically developed in voluntary associations known as standards-development organizations (SDOs), where firms collaborate in developing standards of interest to the industry. As we explain in Section 3, most SDOs in the ICT sector stipulate in their policies that for any (known) patents required to implement their standards, the patent owner must be committed to offering licenses on terms that are “fair, reasonable and non-discriminatory” (FRAND). These FRAND licenses may be royalty-free, as for Bluetooth, HTTP, TCP/IP, and USB, or royalty-bearing, such as for cellular standards, Wi-Fi, and HEVC. In the latter case, FRAND licensing takes on the characteristics of a market transaction.

The market for SEPs and SEP licensing has a number of peculiarities (for a more general discussion, see Charles River Associates, 2016, and Contreras, 2019). First, although standardized products presumably implement a large number of SEPs, licenses are often not sought or finalized until months or years after products are on the market (Contreras, 2013: 59–62). Second, any given mobile telecommunication product is likely covered by thousands of patents declared (potential) SEPs (Baron and Pohlmann, 2018; Bekkers et al., 2020b). Yet due to broad or ambiguous claims and the complexity of the technology, it is often unclear which patents precisely cover a given product, component or standard. Finally, it is often uncertain whether patents declared as “potentially essential” (see Section 3) for a given standard are, in fact, essential, even assuming they are otherwise valid and enforceable. This latter issue is crucial, given that non-essential patents can often be worked around or omitted from a product, unlike essential patents whereby a product by definition has to comply with the respective standard and thus will need to use these essential patents (unless they are only essential to a non-implemented optional feature, or only essential to other device categories). Thus, reliable information about a patent's essentiality goes a long way towards creating transparency regarding its use in a given product.

Yet, as discussed in Section 3.1, studies have reported significant “over-declaration” of SEPs relating to mobile telecommunications (see studies referenced in Contreras, 2019: 211), whether or not intentional. It is believed that many, if not the majority, of the patents declared essential to standards in the mobile telecommunications industry are actually *not* essential. The widespread occurrence of SEP over-declaration creates significant inefficiencies in the market for SEPs. This seems at least partially driven by information asymmetry: owners of (potential) SEPs usually have intimate knowledge about their own patented inventions and whether they are likely to be essential or not. Implementors, on the other hand, are confronted with dozens of SEP holders with thousands of patents, and typically have limited or no knowledge about the details of individual patents claiming to be SEPs. This asymmetry is complicated by long supply chains where products implementing standards range from generic chips, application-specific chips and modules, to intermediate and end products. Downstream firms in these supply chains are often not knowledgeable about the technologies covered by the SEPs in question and implemented in modules they procure from suppliers (see Henkel, 2022).

Asymmetric information and associated uncertainties hamper licensing negotiations for SEPs and invite opportunistic behaviour, creating friction on the MFT and reducing societal welfare. Recognizing these issues, the European Commission (2017: 5) states in its Communication “Setting out the EU approach to Standard Essential Patents”: “*Evidence points to the risk of broad over-declarations and makes a strong case for more reliability with respect to SEP essentiality. Stakeholders report that recorded declarations create a de facto presumption of essentiality in negotiations with licensees. This scenario places a high burden on any willing licensee, especially SMEs and start-ups, to check the essentiality of a large number of SEPs in licensing negotiations. There is therefore a need for a higher degree of scrutiny on essentiality claims.*”

A final inefficiency arising from declarations of patent essentiality is the lack of a consistent method for defining essentiality among SDOs. Bekkers and Updegrove (2013: 35) identify thirteen different features of essentiality definitions in twelve major SDOs (see Contreras, 2007: 12, on additional variants and exclusions). There are two major definitional axes where SDO essentiality definitions differ: (1) the degree to which they cover *optional* portions of a standard, and (2) whether they refer to “*technical*” or “*commercial*” essentiality (commercial meaning that alternatives to the patented solution technically exist but are unattractive or even prohibitive for cost reasons). Yet these terms' definitions are vague (Contreras, 2017: 218), leading to uncertainty among patent holders regarding which patents to declare as essential, and causing implementors to question whether patents declared essential in one SDO would also be essential under another SDO's policies.

These problems are particularly serious given the uptake of IoT applications; compared to smartphone market implementors, IoT implementors are much more numerous and heterogeneous, yet much less knowledgeable about the SEPs that might cover standards used in their IoT devices (Henkel, 2022). Also the growing use and importance of standards to mobility and energy markets, where prospective licensees have less knowledge about patents that may or may not be essential and where different business cultures exist, increase tension and transaction costs. Thus, processes to reliably assess the actual essentiality of declared SEPs are urgently needed to ensure an efficient SEP licensing market.

It is important to acknowledge that full information on essentiality does not remove all possible sources of information asymmetry surrounding essential patents. Other important topics are the validity of patents, the technical merit of the patented technology (impacting the appropriate monetary compensation), and the patents' enforceability (which depends on whether patents are granted, not expired, and their renewal fee being paid, also affected by claims of inequitable conduct, patent misuse, competition violations, etc.). These are all important dimensions in licensing discussions, and for each of them, more transparency may also lower transaction costs. That said, we believe that information on essentiality is the first step in increasing transparency, and other dimensions might come later (once it is determined whether or not a patent is essential in the first place). We note, though, that validity and technical merit are very challenging to assess on a large scale, while other dimensions such as grant and expiry only require looking up a database.[3] Overall, we see essentiality testing as a necessary first step, to be complemented by others, towards reducing information asymmetry and lowering social costs.

3. Determining essentiality in the current environment

Recognizing the nature and consequences of inefficiencies in the markets for SEPs, various studies have attempted to determine the essentiality of patents for given technical standards. While we can learn from these attempts, none to date has provided the market with an assessment of a known accuracy level *and* a cost per patent that would make it feasible to assess all patents potentially essential for a standard.

We first discuss how declarations of potential essentiality are made pursuant to SDO IPR policies (Section 3.1), then recent attempts to assess essentiality using expert-based approaches (Section 3.2), Artificial Intelligence (AI) and other automated approaches (Section 3.3).

3.1. Over-declaration of SEPs

In most SDOs, the disclosure of a patent as potentially "essential" to a standard is based entirely on the patent holder's self-declaration. Accordingly, numerous studies have found significant "over-declaration" of SEPs, particularly at SDOs focusing on mobile telecommunications (Contreras, 2019: 211, collected studies). For example, studies on the GSM (2G) standard and 3GPP standards for 3G and 4G have found over-declaration rates between 8 % and 58 %, with individual patent holders over-declaring at rates as high as 82 % (see Unwired Planet [2017 EWHC 711 at 324–329], citing numerous studies). Such over-declarations may be unintentional or intentional (Contreras, 2017: 223), but whatever the reason, they appear to be endemic within SDOs where patents are required to be disclosed as potentially essential. For a more extensive discussion on over-declaration and its possible causes, we refer to Bekkers and West (2009), Dewatripont and Legros (2013), Contreras (2017), Aoki and Arai (2018), and Lemley and Simcoe (2019).

3.2. Large scale essentiality tests using expert-based approaches

Given the uncertainties and inefficiencies relating to the declaration of SEPs noted above, various mechanisms have been developed to assess the essentiality of sets of patents for a given technical standard. Various parties have conducted such analyses for different purposes. We review these existing approaches and focus on larger scale assessments involving hundreds or at least dozens of patents.[4]

The first patent pools for technology standards, such as those for the MPEG-2 video compression standard and the DVD (digital video disc), appeared in the 1990s (Den Uijl et al., 2013). To ensure their operations were compatible with competition (antitrust) law, some of these pools sought Business Review Letters from the U.S. Department of Justice Antitrust Division. The analysis conducted under these letters concluded that in order to prevent anti-competitive effects, it was important to ensure that patents included in a pool were essential to the standards in question, and that the patented technologies in the pool complemented rather than substituted one another (Gilbert, 2017).[5] As a result, almost all pools set up formal mechanisms for assessing the essentiality of patents proposed for inclusion in the pool. Typically, these assessment procedures: (1) require patent owners to propose patents and submit claim charts that demonstrate why the proposed patents are indeed essential, (2) outsource the assessments to independent, external experts (usually at specialized law firms), and (3) have formal appeals procedures for patent owners and – sometimes – for other pool members and/or licensees. While details of the procedures used by specific pools are usually not made public, an interesting exception is 3G3P, also known as the "WCDMA pool" or "3G Patent Platform", whose initiators published an extensive book describing their approach (Goldstein and Kearsey, 2004). Among other things, this pool (part of the reference set we discuss below) involves independent, parallel assessments and compares their results. A recent study commissioned by the EC includes a review of essentiality assessment mechanisms in patent pools (Bekkers et al., 2020a; see also Merges and Mattioli, 2017).

Pool organizers have extensive experience with such essentiality assessments. Given the strong legal incentives to include only essential patents, pools may be expected to implement diligently performed, high-quality essentiality assessment mechanisms. Also, a pool licensing a portfolio of SEPs for a certain standard should be motivated to build a reputation of licensing actual SEPs; and existing pool members should not be willing to accept a dilution of their portfolio share through newly added patents unless these are actually essential (Merges and Mattioli, 2017; for royalty allocating mechanisms in pools, see Layne-Farrar and Lerner, 2011). The resources they use per patent are in line with such a high-quality assessment.[6] This does not mean pool assessments are perfect. There is inherent uncertainty associated with essentiality assessment. In cases of uncertainty about the actual essentiality of a patent submitted by a member, patent pools may be subject to incentives to include that patent, which would lead to over-inclusion. Also, communication between a submitting member and the evaluator as well as appeal opportunities may lead to over-inclusion. And on an aggregate level, a pool might benefit from more lenient inclusion criteria as it increases its portfolio size compared to other patent holders, thus justifying higher royalty demands.

In addition to patent pools, other parties have assessed the essentiality of patents to different standards. Sometimes, this work has been carried out by academics, who often publish it openly, but more often it is by private consulting firms that only make the results available to parties purchasing their reports. One of the first (published) attempts was by Goodman and Myers (2005), conducted in the context of a conflict between several companies over patent portfolio value (and also sponsored by one of these companies). Many later studies refer to this work, but it was also criticised (see Martin and De Meyer, 2006). From around 2007, a stream of commercial studies followed this approach, including by Fairfield Resources International (Fairfield, 2007, 2009, 2010), which are continuations of the Goodman and Myers (2005)

[3] Validity is difficult to assess since a patent must be compared to the entire state-of-the-art at the time of filing. However, while a considerable share of granted patents may be found invalid if challenged (Henkel and Zischka, 2019, and references therein), errors in granting should be consistent across patent portfolios (and hence less harmful) since they were committed by the patent office not the applicant. Also an assessment of technical merit would be desirable since both SEPs (Bekkers and West, 2009) and patents in general (Scherer and Harhoff, 2000) exhibit a strongly skewed distribution. Yet, technical merit is difficult to assess because a patent must be seen in its interaction with the other elements of the standard and also relative to alternative solutions. Enforceability, finally, depends on a number of factors. Some do not require an assessment (e.g., grant, expiry), while others (inequitable conduct or competition violations) would be very difficult to assess.

[4] We did not review company in-house assessment mechanisms because this information is not shared publicly.

[5] A specific case is where a standard can only be implemented by using one of a set of two (or more) technologies, each patented, making the patents in this set substitutes. In the ETSI definition of essentiality (see Section 4.1), such patents are finally deemed essential, but these are exceptional cases; neither during the 50+ hours of our face-to-face talks with experts regarding their work, nor from the feedback we collected from the assessors in our experiment, was there any indication that this was happening in practice (even though we asked about it explicitly).

[6] These resources are estimated at EUR 5000 to 10,000 for a single European patent, and up to twice as much for a single US patent. The fee per evaluation is typically pre-agreed with the expert performing it and is based on the actual costs and average time spent. (See Merges and Mattioli, 2017 and Bekkers et al., 2020a.)

study, as well as studies from Article One Partners (2011), Cyber Creative Institute (2011), Jefferies and Company (2011), iRunway (2012), and PA Consulting Group (2015). These studies mostly begin with lists of patents declared as potentially essential – for instance patents declared to ETSI – and perform manual assessments of essentiality. It is difficult to assess the quality of these efforts: the underlying methodology, working assumptions and data processing steps are not generally made public, and there is no evaluation, such as a comparison of the results with a benchmark of known accuracy, or inter-rater consistency. Validity of outcomes is not extensively discussed in these works. It is also difficult to compare the outcomes of these studies with each other since they differ in terms of the standard investigated, data selection and cut-off dates.

Finally, large-scale essentiality tests have been performed by economics experts in the context of patent litigation. Such cases include *Unwired Planet v Huawei*, *TCL v Ericsson*, and *In re Innovatio IP Ventures*.[7] The purpose of such analyses was usually to provide input to a court's top-down calculation of FRAND royalties, where the aggregate royalty for all SEPs covering a specific standard is first determined, and then allocated among individual SEPs and SEP holders (Siebrasse and Cotter, 2017). Compared to the analysis by consulting firms discussed above, litigation analyses are somewhat more transparent, yet their procedures vary considerably in design and parameters. While the courts in these three cases recognized that the large-scale essentiality assessments performed by experts were not perfect, they were useful to the courts for their FRAND royalty determinations.

3.3. Large scale essentiality tests using AI and other automated approaches

The use of automated approaches, including Artificial Intelligence (AI), has attracted scholars' attention, and we also see emerging commercial interest in such developments (see IPLytics, 2021). Inspired by work on the computation of semantic similarity between patents (Younge and Kuhn, 2016; Arts et al., 2017), Brachtendorf et al. (2020) investigated the semantic similarity between patents and standards documents to assess actual essentiality of declared SEPs.[8] The algorithm is validated by comparing the findings with the results of the manual essentiality assessments for the TCL v Ericsson court case mentioned above. At the individual patent level, the Brachtendorf study found limited consistency between its outcomes and the court case data. From the set of 166 patents assessed essential by manual evaluators, the automated system predicted only 40 (24 %) were essential. From the set of 236 patents assessed not essential by manual evaluators, the automated system predicted 216 (92 %) were not essential. If we assume the reference point is perfect (which it may not be), then the automated system has many false negatives, and fewer false positives. Yet the authors found strong and highly significant correlations between the experts' decisions on standard essentiality and their measurement of semantic similarity. All in all, these initial results are promising but not yet satisfactory in terms of predicting essentiality on an individual patent basis.

An undeniable strength of automated approaches is their scalability. Potentially, they allow the analysis of very large sets in a relatively short time span, and at low costs. Also, AI systems could provide objectivity beyond that of human assessors who may be prone to certain political or economic motives (less of a risk in our experiment as it was conducted by patent examiners, whom we believe to be impartial). But automated systems also come with inherent limitations. Firstly, the meaning, interpretation, and precise scope of words and terminology (both in patents and standards) depend on the context, making it hard to automate. Second, semantic approaches can face difficulties dealing with changes in terminology over time. Third, the patent to be evaluated, or parts of it, may be written in a different (natural) language than the respective part of the standard. Furthermore, even with the same natural language, the vocabulary in patents (drafted by patent attorneys) often differs from standards (drafted by engineers). Fourth, a technology or solution required to implement the standard may not be explicitly mentioned in the standard's text but may still be required in order to satisfy the standard (i.e., be implied by the standard). Fifth, an essentiality analysis should consider all existing (patented[9] and unpatented) technical alternatives that may also satisfy the standard and render the patent under investigation non-essential. For instance, when a standard requires a quasi-random code for some function, a certain patented technology may indeed be able to generate such a code, but the existence of other solutions that also generate a code which would satisfy the standard's requirements must also be considered. An experienced human assessor, well-trained in the technical field, would be expected to have such knowledge. Current AI systems in this field, however, merely compare a focal patent and a focal standard. Adding the whole universe of external solutions would be challenging given the current state of relevant AI technology. We acknowledge that several of the above aspects (especially different language spheres and alternative technological solutions) can be challenging to human assessors as well. At the same time, we believe that human assessors, being aware of the issues and addressing them properly, can do a better job. In fact, dealing with such challenges is already inherent in a patent examiner's regular tasks.

In addition, there are institutional limitations. Automated systems (as currently explored in this context) are not transparent and explainable, and consequently lack accountability. This may seriously hamper stakeholders' acceptance of the system's output. Furthermore, any automated system is prone to gaming, whereby patent owners, anticipating the workings of such a system, adapt the wording of their patent applications (which might end up in the granted patent claims) to the wording of standard documents.[10] Such gaming challenges the objectivity of the automated system.

In sum, these approaches to essentiality assessment provide useful input in terms of designing an essentiality testing mechanism, but do not yet answer the question of whether essentiality assessments can be made sufficiently efficient (in terms of time and costs) as well as sufficiently accurate, to set up a large-scale system and thus overcome important inefficiencies in the market for SEP licensing. At the same time, it may be feasible to create automated systems that perform pre-screening (i.e.,

[7] Unwired Planet v Huawei, [2017] EWHC 711 (Pat); TCL Commc'n Tech. Holdings, Ltd. v. Telefonaktiebolaget LM Ericsson, No. 8:14-cv-00341 (C.D. Cal. Dec. 21, 2017), rev'd on other grounds, 943 F.3d 1360 (Fed. Cir. 2019); Innovatio IP Ventures, LLC Patent Litigation, 2013 U.S. Dist. LEXIS 144061 (N.D. Ill. Sept 27, 2013).

[8] The authors identified standards documents based on patent declarations at ETSI, resulting in 4796 standards documents, and compared them with 37 million patent documents, considering patent claims as well as technological descriptions. The study used an algorithm developed by Natterer (2016).

[9] Whether the existence of patented alternatives removes essentiality depends on the exact definition of essentiality adopted by the SDO in question. The definition of essentiality at ETSI is explicit on this aspect: if alternatives exist that are not patented, the patent in question is not essential; if only alternatives exist that are also patented, then the focal patent is essential (as well as the patented alternatives). (ETSI, 2021: Annex 6, §15, Item 6; see also Contreras, 2017: 218). Rules at other SDOs may differ or are not explicit (Bekkers and Updegrove, 2013: 66).

[10] For example, the U.S. Federal Trade Commission found that semiconductor designer Rambus, Inc. deliberately modified the claims of its patent applications during prosecution better to cover technology features concurrently being incorporated in JEDEC standards. In the Matter of Rambus, Inc., FTC Docket No. 9302, 2006 FTC LEXIS 101, *6 (FTC, Aug. 20, 2006) ("through its participation in JEDEC, Rambus gained information about the pending standard, and then amended its patent applications to ensure that subsequently-issued patents would cover the ultimate standard"); see also id. at *88–89 (describing meetings between Rambus CEO and its JEDEC representative concerning "how Rambus might add claims to cover JEDEC standards").

systems with very few false negatives that allow a relative degree of false positives) or automated systems that would otherwise assist manual assessments.

4. Experiment design and data

4.1. Definition of essentiality

SDOs and other organizations have adopted different definitions of essentiality (see Section 2.2). Since our experiment considered assessments of patent essentiality for ETSI standards, we adhered to that SDO's definition of essentiality: "'*ESSENTIAL' as applied to IPR means that it is not possible on technical* (*but not commercial*) *grounds*, *taking into account normal technical practice and the state of the art generally available at the time of standardisation*, *to make*, *sell*, *lease*, *otherwise dispose of*, *repair*, *use or operate EQUIPMENT or METHODS which comply with a STANDARD without infringing that IPR*. *For the avoidance of doubt in exceptional cases where a STANDARD can only be implemented by technical solutions*, *all of which are infringements of IPRs*, *all such IPRs shall be considered ESSENTIAL*." (ETSI, 2021).[11]

While preparing the experiment and discussing this definition with the various patent offices involved, some offices raised concerns that patent examiners were not trained in determining infringement and asked if the assignment could be re-phrased. Together with these offices, we developed an alternative that we call a "novelty-based test", based on the following: *in the hypothetical case that the standard document had already been published before the patent priority date*, *would that document have been novelty-destroying*? Several stakeholders indicated that they did not expect the specific definition to make a difference in essentiality assessments; nonetheless, in our analysis we compared the essentiality assessments based on the ETSI definition to those based on the novelty-based test.

4.2. Source and characteristics of the patents to be assessed

In order to interpret the outcome of the experiment, it is important to understand the features of the set of patents that will be assessed if our approach should be used in a large-scale assessment system. We assume that these patents will be selected for assessment by their respective owners. Compared to other selection scenarios (which were found to be less realistic in terms of organizational feasibility by Bekkers et al. (2020a), such patents are the most challenging to assess since they were selected by their owners as being presumably essential. This choice implies that expected accuracy will be lower than if we had opted for a different scenario. Therefore, our outcomes in terms of accuracy constitute a lower bound for any of the designs.

4.3. Reference assessments

To determine the accuracy of the assessments in this experiment, a reference was required for comparing our results. The ultimate, authorised decision concerning essentiality lies with competent courts. While some courts have indeed issued (public) verdicts on the essentiality of patents, the number of data points is very limited. Moreover, there is the risk that our assessors were aware of these court verdicts, thereby creating a possible bias.

For our experiment, we used what we believe are the most accurate assessment points existing *outside* a litigation context: the assessment by patent pools. To comply with competition (antitrust) law, these pools have developed diligent and sophisticated procedures where patents submitted to the pool are scrutinised by external, independent parties (usually law firms or patent attorneys specialized in this task), as discussed in Section 3.2. While pool assessments cannot be regarded as perfect, they are considered by almost all stakeholders as the gold standard, and we believe they are appropriate as a reference point for our study.

4.4. Selection of assessment cases and associated data collection

To perform the experiment, we developed a sample of cases, where "case" refers to a combination of a granted patent document and a (specific release of a) standard document (e.g., TS 25.211 V2.5.0). While "positive" reference cases can be easily identified using public information by pools on which patents were determined essential, "negative" reference cases required a different approach. Ideally, we wanted to know which patents were actually submitted to pools, but then rejected. To obtain such information, we sought collaboration with patent holders, and, after negotiations, several patent holders participating in pools were willing to share that (private) information. Moreover, we also found them willing to share the claim charts that they actually submitted to the pools, for both accepted and rejected patents, allowing us to provide exactly the same information to our assessors as provided to the pools. The highly confidential nature of these claim charts required non-disclosure agreements.

To ensure our assessors had in-depth expertise in the field of the standards and patents they were reviewing, we focused the experiment on a single technological area, namely ETSI/3GPP 3G and 4G standards, and selected assessors accordingly. This technical area is one of those primarily calling for essentiality testing (European Commission, 2017). Furthermore, there are several patent pools active in this area, allowing us to use their essentiality decisions as reference points.

Our experiment used data from the following pools that all have a licensing program for these patents: the "WCDMA" patent pool, the Sisvel LTE/LTE-A patent pool, the Via Licensing LTE patent pool, and the Avanci patent pool. While there is no strong reason to believe that the difficulty in assessing a patents' essentiality depends on the specific company or where it is situated in the value chain, we nevertheless wanted to have a mix. The above pools on which we based our data have licensors from different positions in the value chain: combined, 37 % of the licensors are vertically integrated firms, 36 % are upstream technology firms, and 27 % are telecommunications network operators. This mix is also reflected in the patents selected for the experiment.

Given that we engaged patent examiners from European patent offices, we only included EPO patents in the experiment. While granted EPO patents always include patent claims in the English language, the other text in the document may be in German and French, also official EPO languages. To prevent possible bias in patent selection by language exclusion, and to be able to discover whether language issues would affect the assessment, we decided not to remove patents where parts of the patent text were in German or French (but the claims were stated in English).

Our final data set has four categories, as shown in Table 1. Categories I and II are based on the data provided by patent owners, discussed above. There are two pertinent points concerning Category II. Firstly, because companies usually review their patents internally and only submit patents to pools which they believe likely to succeed, this set is smaller than Category I. Secondly, because of this preselection, these patents may be relatively more difficult to assess. We will return to this in our analysis and conclusion. Categories III and IV are based on public data, and complement the above data to ensure we had the required number of patents for our factorial design (see below). Category III are patents publicly disclosed by the pool as essential. We had no claim charts for these, so we used them for the cases where we did not plan to provide our assessors with claim charts. Category IV is the most challenging because patent pools do not publish the identity of patents that were submitted but rejected. We had to reconstruct this category, by creating a set of patents very similar to the one in the actual pool, using a

[11] While the reference is to the most recent version of the ETSI policy, there are no substantive differences relevant to our paper with the policy in place at the time of our experiment.

Table 1
Data sources.

	Patent included in pool	Patent not included in pool
Based on data provided by patent owner	Category I	Category II
	- Data sources: pool acceptance information supplied by patent owner (verified by public pool information); claim charts supplied by patent owner - Data quality: very high - Assessment difficulty: average	- Data sources: pool rejection information and claim charts supplied by patent owner - Data quality: very high - Assessment difficulty: high
Based on public data	Category III	Category IV
	- Data source: pool inclusion information from pool publication - Data quality: very high - Assessment difficulty: average	- Data source: pool non-inclusion data reconstructed - Data quality: medium - Assessment difficulty: average

series of defined criteria.[12] By having to reconstruct this set, we acknowledge that the data quality may be lower than in the other three categories. One possible issue is that the data for identifying the proper standard document for a patent (based on potentially essential declarations to ETSI) may be less reliable than for the other three categories (where that data is based on pool data). As discussed in Section 4.6, we paid special attention to any signs of such issues, and where necessary discarded observations.

In total, the experiment involved 43 unique patents (not members of the same INPADOC family) and 48 unique standards documents.[13] These patents are owned by vertically integrated companies (36 %), upstream technology firms (53 %) and network operators (11 %). Table 2 provides additional descriptive information about the patents used. Cases were randomly allocated to assessors, ensuring that an individual assessor did not receive more than one case on the same patent or on the same standard document (to prevent unobserved learning effects).

Table 2
Descriptive data on the patents in the experiment.

	Average	St. Dev.	Min.	Max.
Priority year	2001	4.16	1993	2008
Application year	2002	4.64	1994	2015
Grant year	2007	4.40	1997	2017
Grant lag (months)	58.05	27.05	28	137
No. of claims in granted patent	17.12	9.35	1	44
INPADOC family size	14.49	7.91	3	29
DOCDB family size	12.65	8.01	3	29

The experiment follows a factorial design. The factor we are mainly interested in, the treatment, is whether a patent has claim charts or not. Blocks include the reference point, the availability of claim charts, and the definition of essentiality used in the assessment (ETSI vs. novelty-based). Uncontrollable nuisances are accounted for by randomizing.

4.5. Assessors, case allocation, and instructions

Twenty patent examiners employed by six different European patent offices participated as assessors, selected on the basis of their considerable expertise in the technical areas of our cases. During a debriefing survey, we collected self-reported information from the assessors relevant in terms of their ability to carry out the task. Descriptive statistics are shown in Table 3 (of the 20 assessors, 18 fully provided this information). The results confirm that the assessors deem themselves highly knowledgeable in the technical areas relevant to our cases, are familiar with the type of documents they had to handle except claim charts prepared by the actual patent owner (67 % of respondents report they use them “frequently” or “all the time” but the remaining 33 % were less familiar with them). We do not see that as a problem, however, for answering our related research question Q2. Finally, they have high language proficiency in English. Most assessors report CEFR B1 level language proficiency in German and French – an aspect we will return to later.

While the management of patent offices was closely involved in the study design and operationalisation, none of the assessors was given any of this information. During debriefing, this was again confirmed. Each assessor received eight different cases, equally distributed in the combination of essential/non-essential and with or without claim charts, yielding 40 cases per combination (see Table 4).

The assessors were asked to carry out their tasks during regular work time, and their employer indicated their contribution to this project was much appreciated. Assessors did not receive direct (financial) compensation for their work. We believe patent examiners have a very strong work ethic and moral standards about the quality of their work, also reflected in our communication with them. Assessors were asked to spend as much time on a patent as needed until they felt sufficiently confident about their decision or inform us if they felt they could not further increase their confidence. After the experiment was over, the assessors reported they spent on average 7.7 h per patent, but we note that for 24 % of the cases, assessors took 16 or more hours, confirming they did take the time when they deemed it necessary (see also Section 5.4). For 11 % of all cases, assessors reported they did not feel sufficiently confident about their results.

Assessors were instructed not to assume any particular distribution of cases, and actually did not know we were using a reference point.

Each case, as indicated above, refers to a combination of a granted patent document and a (specific release of a) standard document (e.g., TS 25.211 V2.5.0). Standards, and patents essential to them, have an *n:m* relationship, and also in our dataset, some cases shared a patent or standard. To avoid bias from learning, individual assessors considered every patent and every standard only once throughout the entire experiment. With the above restrictions in mind, both the allocation of the cases and the order in which the assessor processed them were randomised (and our logistics ensured they were indeed evaluated in that order).

Assessors were given an extensive set of instructions, developed together with the patent office management departments, and pre-tested for clarity (see Appendix A: “Instructions for assessors”). Assessors were instructed not to search for any information in addition to what we provided so that their assessment was solely based on the patent text and the standard document provided. They were not allowed to discuss cases with others for the entire duration of the experiment. Also, they were

[12] These criteria were: (1) the patent owner is a member of the WCDMA pool, (2) the patent was declared to ETSI as potentially essential for the relevant standards, (3) the ETSI declaration included information on the specific standards documents for which the standard was potentially essential, (4) the ETSI declaration was within a time window when the declaring firm declared most of its patents that eventually became WCDMA pool patents, (5) the patent is not part of WCDMA pool patents nor of an INPADOC family containing other patents among WCDMA pool patents, and (6) the patent was applied for at the EPO and granted.

[13] The following standards document were assigned the most (20 times or more): TS36.211, TS25.212, TS25.211, TS31.102 and TS25.223, although individual cases could refer to different release dates/document versions. (e.g., TS 25.211 V2.5.0) to ensure they were identical to the test in the reference. Note that for each case, the relevant version of the standards document was selected and given to the assessor; 27 % of the cases came with a single standards document; 32 % with two, 16 % with three and the remaining 2 % with four or five.

Table 3
Assessors' descriptive statistics.

		Average	Median	St. Dev.	Min.	Max.	Obs.
Self-reported technical expertise in specific areas (1)	Radio layer communications protocols	4.83	5	1.07	2	6	18
	Physical layer radio protocols	4.67	5	1.15	2	6	18
	The UMTS (3G) standard	4.44	5	1.17	2	6	18
	The LTE (4G) standard	4.56	5	1.26	2	6	18
Self-reported familiarity with specific document types (2)	Standards documents	3.17	3	0.76	2	4	18
	3GPP standards documents	3.17	3	0.76	2	4	18
	Patents	3.94	4	0.23	3	4	18
	Telecommunications technology patents	3.67	4	0.58	2	4	18
	Claim charts (made by the patent owner)	2.72	3	1.19	1	4	18
Self-reported language proficiency (3)	English	6.50	7	0.69	5	7	18
	German	4.00	4	2.38	1	7	18
	French	4.89	5	1.63	2	7	18

(1) 1="No knowledge", 2="Knowledge of existence", 3="Understanding of concept", 4="Basic understanding", 5="Professional understanding", 6="Expert knowledge".
(2) Scale is: 1="Never", 2="Once or twice", 3="Frequently", 4="All the time/professionally".
(3) The scale is based on the Common European Framework of Reference for Languages (CEFR) plus one flanking category 1="No understanding", 2 = A1 level, 3 = A2 level, 4 = B1 level, 5 = B2 level, 6 = C1 level, 7 = C2 level.

Table 4
Allocation of assessments.

		Block		
	Claim chart	Essential	Non-essential	Total
Treatment	With	40	40	80
	Without	40	40	80
Total		80	80	160

only allowed to look up *technical* information from other sources if this was necessary to understand the technology described in the patent or standard (e.g., a technical handbook or standards document in the same 3GPP series). Patent documents were anonymised by removing patent number and assignee information, and assessors were instructed not to look up information on the specific patent (e.g., by searching for the title). For additional verification, assessors were asked to indicate if they recognized the identity of the patent owner and/or the patent. Moreover, assessors were instructed to consider all claims in the patent but could stop their assessment once they found one claim that they considered essential. We also gathered qualitative outcomes by asking the assessors both closed questions and open feedback. Finally, after the experiment was completed, all assessors received a feedback and debriefing form.

4.6. Data verification

Before carrying out the data analyses, we verified the assessment data for factors that could potentially have a confounding effect on the experiment. During debriefings, we learned that one group of participants had not respected all the elements in the instructions. Despite their good intentions, they did not realise this was at odds with our research design. We had to exclude the associated observations from our quantitative analysis, but we still used their feedback in the qualitative analysis. Furthermore, there were a few instances where: (a) participants reported they had seen the patent before (possibly as an examiner), (b) participants informed us they knew (or thought they knew) who the patent owner was, (c) participants reported specific issues; for instance, participants doubted whether essentiality may lie in another standards document than the one we provided, (d) there were two or more assessments for a case and at least 75 % disagreed with the reference point. In total, these affected five cases (two in Category II and three in Category IV) associated with 19 observations. To ensure high data quality, we discarded these 19 observations from our analysis.[14] Our final analysis included 109 valid observations.

5. Results

To present the results of our experiment, we address the overall accuracy level, then the impact of claim chart availability on accuracy, and finally, the differences between the original and the alternative (novelty-based) definition of essentiality. For each, we discuss both the quantitative and qualitative outcomes.

5.1. Quantitative findings on the overall accuracy of assessments

Table 5 shows the assessment outcomes compared to the reference points, in a confusion matrix. The assessment achieves an accuracy of 74 %, i.e., in 74 % of the cases, the outcome is consistent with the reference. The sensitivity of the assessment (the share of consistently classified patents among those found essential by the pool) is 83 % and thus higher than its specificity of 62 % (the share of consistently classified patents among those found non-essential by the pool). This difference, should it persist in a large-scale implementation of the assessment procedure, becomes relevant the more the initial set of patents is imbalanced, i.e., contains unequal shares of actually essential and non-essential patents. It also becomes relevant if correctly identifying actually essential patents (i.e., avoiding false negatives) should be more important than avoiding false positives, in which case a high sensitivity is more desirable than high selectivity. The precision (share of patents found essential by the

Table 5
Discrimination between essential and non-essential patents.

	Experiment outcome		
Essentiality status according to the reference point	"Essential"	"Non-essential"	Total
Essential	53 (83 %)	11 (17 %)	64 (100 %)
Non-essential	17 (38 %)	28 (62 %)	45 (100 %)
Total	70 (64 %)	39 (36 %)	109 (100 %)

Note: Cells show the number of observations and percentage of row total.

[14] We performed all the statistical tests separately as presented in our analysis below, while including these 19 observations, but this did not change the significance of our findings.

pool among those assessed as essential), finally, is 76 %. A chi-square test confirms, unsurprisingly, that the assessors are significantly better than random in differentiating essential from non-essential patents ($p < 0.001$).

5.2. Quantitative findings on impact of claim charts and essentiality definition of accuracy

Table 6 shows the result of the assessments depending on the availability of claim charts (52 of the 109 observations included a claim chart). The percentage of assessments inconsistent with the reference was half as large with claim charts (17 %) than without claim charts (33 %),[15] which constitutes a significant improvement ($\chi^2 = 3.66$, $p = 0.056$). In a separate test, we investigated whether claim charts affect Type I and Type II errors differently, but we did not find a statistically significant effect here.

As mentioned in Section 4.1, most of the assessors applied the ETSI definition of essentiality, while others kept to what we call a novelty-based test. So far, the results shown included the data points from both. To check this was indeed valid, we verified whether the results of the "ETSI-based" essentiality assessments and those based on the novelty-based tests are comparable.[16] Table 7 shows the results.

The difference between the distributions in Table 7 is insignificant ($\chi^2 = 0.358$, $p = 0.55$). If anything, we see that the outcome of novelty-based assessments is slightly more often consistent with the reference (79 %) than the regular assessments (73 %). This result has an important implication. Even though most assessors felt qualified to perform the assessments (expressed in their open feedback), patent examiners are not always trained to perform infringement analyses, and infringement partly depends on the respective national law. However, since patent examiners are trained to perform novelty analyses, they can directly perform essentiality assessments under the novelty-based definition. This is relevant given that stakeholders expressed their confidence in the reputation of patent offices as trustworthy, independent third parties, thus qualified candidates to perform these tests on a large scale.

5.3. False positives vs. false negatives

An intriguing aspect of our experiment is the comparison between false positives and false negatives. As discussed in Section 5.1, the sensitivity of the assessment (83 %) differs from its selectivity (62 %), in other words, the share of false negatives is smaller than the share of false positives (17 % vs. 38 %). Comparing this result with an expected outcome that would give the same likelihood of false positives and false negatives shows significant differences ($\chi^2 = 5.87$, $p = 0.015$).

Table 6
The effect of claim chart availability.

	Experiment outcome compared to reference point		
Claim chart availability	Consistent	Inconsistent	Total
No claim chart	38 (67 %)	19 (33 %)	57 (100 %)
Claim chart	43 (83 %)	9 (17 %)	52 (100 %)
Total	81 (74 %)	28 (26 %)	109 (100 %)

Note: Cells show the number of observations and percentage of row total.

Table 7
Novelty-based vs. regular essentiality definition.

	Experiment outcome compared to reference point		
Type of essentiality definition	Consistent	Inconsistent	Total
Regular	59 (73 %)	22 (27 %)	81 (100 %)
Novelty-based	22 (79 %)	6 (21 %)	28 (100 %)
Total	81 (74 %)	28 (26 %)	109 (100 %)

Note: Cells show the number of observations and percentage of row total.

There are two possible interpretations of this result. Assessors might be more inclined to assess a patent as essential than non-essential, as a form of confirmation bias. Or, our non-essential cases (according to the reference) might have been more difficult to assess than the essential cases. In their qualitative feedback, the assessors in fact indicated that proving a patent's essentiality typically took less effort than proving non-essentiality, particularly when no claim charts were provided.

We also explored whether claim charts impact false positives and false negatives in different ways; one might hypothesize, for instance, that a claim chart for an actually essential patent simplifies the assessor's task of establishing essentiality, while a claim chart for a patent that is actually not essential makes it even harder to establish non-essentiality. Table 8 shows the data, but the difference between the effects fails to be significant, according to a Tarone-adjusted Breslow-Day test ($\chi^2 = 2.35$, $p = 0.125$), which might be due to the number of observations we have.

5.4. Qualitative findings

We also gathered extensive qualitative feedback from the assessors, both for each individual assessment they performed, and after having finalized all their work. They often ventured to offer us very elaborate and detailed information on how the assessment took place, the challenges they encountered, their approach, and their understanding of the patent. This information helped us to assess whether there were elements that affected the outcomes that we had not anticipated or had overlooked. Summarizing, they expressed their strong belief that the task they were given required a thorough knowledge of the standard documents. Such knowledge, they felt, could be gained in practice by specialising in essentiality assessments. For example, in one case, the patent referred to a feature that was not available at all in the specific standards document given to the assessor but may have been elsewhere in the standard (3GPP standards cover hundreds of separate documents). In another case, an assessor suspected that a patent would be essential for a newer release of the relevant standards document, whereas the version provided did not require the use of the patented technology. Assessors indicated that improved searching tools could help a less experienced assessor, especially where the claimed essential features were spread over a combination of standard documents. Moreover, assessors felt they would have benefited from access to additional information about the patent, such as written opinions from patent offices, claim trees, external knowledge and interaction with stakeholders. Additionally, while the text in the standard (and, where applicable, the claim charts) we provided was always in English, and the claims in all

Table 8
Separated results for false positives and false negatives depending on claim chart availability.

		Experiment outcome compared to reference point		
Claim chart provided	Reference point	Consistent	Inconsistent	Total
No	Essential	23	9	32
No	Non-essential	15	10	25
Yes	Essential	30	2	32
Yes	Non-essential	13	7	20
All		81	28	109

[15] Note again that our non-essential cases with claim charts might have been the most difficult to assess, since patent holders would only create a claim chart if they believed the patent was essential. Thus, for the overall population of patents potentially candidates for an assessment procedure, the difference might be greater than Table 6 indicates.

[16] Note that we ran this analysis *before* the others, but only discuss this analysis now for readability reasons.

EPO patent publications are in English, we had cases where other part(s) of the remaining patent text were in German or French. While the scope of the patent is defined by the claims, the other text may be useful to interpret those terms. Around 60 % of assessors reported having at least a B1 language proficiency level in German and French, which also means that some 40 % had a lower proficiency level. Overall, two out of 20 assessors provided feedback that they felt language issues could lead to difficulties in making the correct assessment. Implementing the above suggestions should help to increase the percentage of consistent assessments.

Assessors reported spending on average 7.7 h per case, but as expected, the times differed considerably, because some cases could be (perceived as) more difficult than others. For 24 % of the cases, assessors took 16 h or more (Fig. 1). For nine observations, participants reported spending "much more" time than anticipated; for five observations "much less"; and for other observations, only a little more or less than expected. Using the hourly rate currently applied by commercial patent attorneys working for patent pools,[17] this translates to a cost of approximately EUR 2400 to 3200 per assessment. Yet, the assessments proposed in this paper were performed by patent examiners, not attorneys, and costs/fees might be quite different.

Participants appeared relatively confident in their evaluations, labelling them as "very certain" (25 observations) or "quite certain" (101). In the remaining observations, participants felt "undecided" (16), "quite uncertain" (12) or "very uncertain" (6). In Fig. 2, we show the levels of self-reported confidence in relation to the assessment outcome and the consistency (i.e., outcome is accurate or not). The level of confidence is not associated with the consistency of the evaluation ($\chi^2 = 1.01$, $p = 0.29$). Yet, the self-reported confidence level is significantly higher (ordinal regression; $p = 0.002$) for "essential" evaluations than for "non-essential" evaluations. Participants generally (in qualitative feedback) indicated they felt qualified to perform the assessment, even though it was a new task, and a few cases were noted as outside the assessor's regular field of expertise (at a lower-level technical "layer"). In 137 cases, participants reported that their skill level increased "slightly" or more, which may indicate the presence of a learning effect. See Bekkers et al. (2020a) which elaborates on the learning effect observed in the earlier experiment (Fig. 2).

Participants indicated (on a closed question form) that the claim chart was "very helpful" (31 observations) or even "extremely helpful" (28 observations).[18] In their open feedback, assessors indicated that claim charts were useful for two reasons: they saved time and made them feel more confident about the outcome. They also commented that in the absence of claim charts, the process of reading the patent description, isolating the parts truly reflected in the claims, and then doing the same for the standard document and matching both parts, required a lot more effort.

6. Conclusion and discussion

The purpose of this study was to investigate whether essentiality assessments can be made sufficiently efficient (in terms of time and costs) as well as sufficiently accurate to set up a large-scale system of essentiality assessment, and thus overcome important inefficiencies in the market for SEP licensing.

In our experiment, where assessors spent an average of 7.7 h on each case, we found that 74 % of the outcomes are consistent with patent pools' essentiality assessments, which we used as reference point. When the assessors were provided with claim charts as additional input and used the "regular" essentiality definition, consistency increased to 84 %. Given our decision to study the organisationally most feasible, but in terms of accuracy most challenging assessment scenario (see Section 4.2), we believe these results are encouraging, and can be considered as a lower bound for other designs likely to have much higher accuracy scores.

Further improvements should be achievable when implementing our approach in practice. Firstly, we introduced several limitations to ensure a proper research design. Participating assessors were not allowed to work in teams or exchange information, nor look up patent prosecution history. Apart from the information we provided, the only other source they could look up was general technical background information such as the meaning of a specific term. In a practical implementation of the approach, such limitations would not be imposed, thus likely improving performance. Second, even though the assessors were selected based on their expertise in the relevant ETSI/3GPP 3G and 4G standards, this is still a relatively broad area. In practice, a larger assessment team would include specialists in relevant subfields (switching, radio protocols, etc.), and patents could be allocated to assessors according to their key technological competences, thus improving performance. Thirdly, our data set included a significant number of relatively "difficult" cases, such as patents previously submitted to a pool but subsequently rejected. Patent pool assessments are costly, and we must assume that the patent owner had reasonable expectations that the patent was essential. By contrast, in a large-scale implementation of our approach, assessments may start from *all* patents declared as potentially essential to SDOs, implying that there will be many more obvious (i.e., "easy" to assess) cases. Again, performance should improve. Fourth, we see many learning opportunities, both for individuals (progressing experience and knowledge) and in a team setting. Indeed, the participating assessors commented that the eight cases were not enough to generate learning effects. Also, if assessments are carried out on a larger scale, one can ensure assessors are only offered documents in a language they master at a high level, thus further improving the system's performance. Finally, a system implemented in practice could allow the patent owner involved as well as third parties (such as implementers) to appeal the result of the assessment. This measure should improve the accuracy of the assessments, though we recognize such procedures need to be designed carefully to avoid potential misuse.

Given the outcome of our experiment and the above opportunities to improve performance, we believe that a large-scale system of essentiality assessments based on our approach can achieve a good degree of accuracy at a comparatively low cost (compared to the currently known processes with high accuracy) and overcome important inefficiencies in the market for SEP licensing.

We note that a large-scale assessment system, as studied in this paper, is not the only way to create more transparency in terms of essentiality. An alternative approach is a two-stage disclosure procedure within SDOs, where companies at an early stage disclose patents that are potentially essential. Later, once the standard is frozen and the potentially essential patent is granted, companies re-assess their patent and make an additional disclosure whether or not they believe the patent is actually essential. While such two-stage procedures have been proposed (Qualcomm, 2006), no SDO has reached a consensus to include such an obligation in its IPR policy, and given the SDO governance structure (see Baron et al., 2019), it is unlikely this will happen any time soon. Moreover, such a two-stage procedure will still be a self-declaration by the patent owner, not an assessment by an independent, impartial third party that applies consistent standards.

Our main contribution in this paper is empirical and policy related. Complementing this, we provide a theoretical contribution concerning the notice function of patents. Property rights and the ability to enforce them are central to the functioning of markets. A key characteristic of property rights in this regard is their notice function: they should unambiguously signal if an activity infringes on a right. As Bessen and

[17] Bekkers et al. (2020a), find pools spend between EUR 5000 and 10,000 per evaluation, taking 2 to 3 days on average, i.e., average fees are EUR 312 to 416 per hour. The pool assessments are performed by attorneys, chosen to match the patent's applicable jurisdiction. Their fees differ greatly between jurisdictions.

[18] These numbers include qualitative feedback on cases we had to exclude from the quantitative analysis.

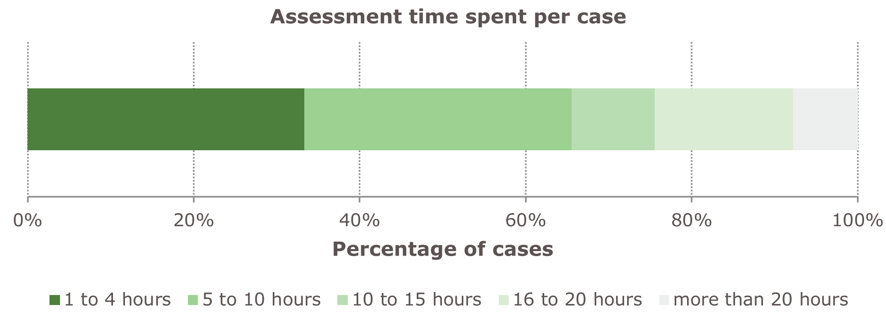

Fig. 1. Time spent on assessing each case.

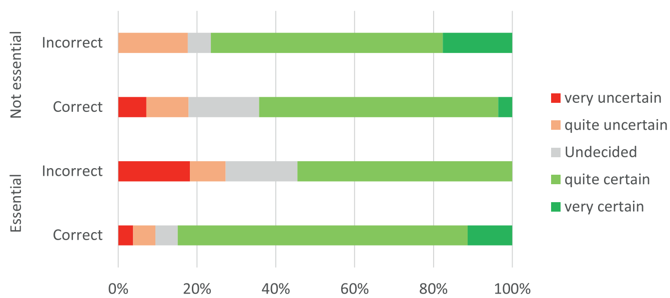

Fig. 2. Assessors' self-reported confidence by outcome and accuracy.

Meurer (2009: 8–9) point out, patents' notice function is frequently ineffective. This is for example because patents often have unclear boundaries due to abstract language that requires interpretation, and many patents might claim aspects or features of any given product. For all practical purposes, the uncertainty that comes with an imperfect notice function relates not just to the patent, but also to the product that potentially practices the patented invention. The problem of determining whether this is indeed the case usually involves the patent owner and the maker of the artefact. For a patent not related to any standard, improving its notice function with respect to the artefact at hand is a private good of the two parties involved. The situation is different when it comes to standards. Open technical standards such as LTE are implemented by numerous actors in a multitude of devices. Thus, the question of whether a product implementing the LTE standard infringes a specific patent concerns all implementers in the same way (apart from patents on optional features of the standard). Clarity regarding which patents are actually practised by LTE devices would benefit all implementers, but also patent owners since it might affect their share of aggregate royalties. Thus, improving patents' notice function on open standards has the characteristics of a public good, a theoretical insight that justifies public ordering.

Our experiment has several limitations: (1) We used patents declared at ETSI as cases and the ETSI essentiality definition. Performing such assessments using other SDOs' definitions of essentiality could be more challenging, especially if they are based, for instance, on "commercial" rather than technical essentiality. (2) The patents in our sample are subject to two selection effects: they are owned by companies that participate in pools (all patents) and the patents themselves are factually offered to a pool (Cat. I, II and III). Some studies discuss possible bias in such selections. Layne-Farrar and Lerner (2011) find that vertically integrated firms (that both perform R&D and manufacture products) are more likely to join a pool, in line with a model analysis by Aoki and Nagaoka (2005). Layne-Farrar and Lerner furthermore argue that firms with higher value patent portfolios are less likely to join certain types of pools. Yet, we observe that the pools relevant to our data are actually very diverse in membership across business models (see Section 4.4). And even if pool participation correlated with patent "quality" (e.g., technical merit, value, legal quality), we would not expect such value to be associated with the difficulty of assessing essentiality. Concerning the selection of patents by pool participants to be offered to pools: participants have a strong incentive to offer all patents which they think may be essential because it increases their share in the pool royalty distribution. Interviews with all pools and pool participants confirmed that this is indeed true for the pools and pool participants in our dataset. Thus, we are confident that these selection processes do not affect our results in other ways than already discussed. (3) While ETSI requires parties to disclose which specific patents they believe are potentially essential, other SDOs, including ITU, IEEE, and ISO/IEC, allow parties to submit "blanket" declarations that do not indicate specific patents. While a large-scale essentiality test mechanism does not necessarily rely on declarations made at SDOs (it may also start with patent owners proposing their patents for assessment), this may limit system design options. (4) The availability of input claim charts, where we observed the highest degree of consistency, will depend on whether patent owners are willing to make such information available – and, in turn, what will motivate patent owners to do so. As indicated above, such questions concerning "institutional feasibility" are beyond the scope of this paper but are addressed in complementary work presented in Bekkers et al. (2020a). (5) Differing patent prosecution procedures, examination details, linguistic translation variations, third party interventions and substantive legal rules across jurisdictions may lead to differences in the scope of patent claims issued from one jurisdiction to another, even for patents in the same family and originating from the same international (Patent Cooperation Treaty) application. Determining, even with a high degree of certainty, that a particular member of a patent family (e.g., a European patent) is essential to a given standard may only provide approximate information about patents issued in other countries, for example in the U.S.A., China or Japan. (Note that EPO patents, however, are identical across all countries for which the patent owner chose national validation.) (6) While we believe that patent pool assessments are an appropriate reference for this study, they do not represent an absolute reference point, and as mentioned above, such a reference point does not exist. Therefore, our findings are necessarily limited to observing consistency, not accuracy.

There are ample opportunities for future research in this area, especially since the European Commission announced in its November 2020 IP Action Plan that it will explore the creation of an independent system of third-party essentiality checks with a view to improving legal certainty and reducing litigation (European Commission, 2020a: 13). One of these opportunities is research on understanding how AI-based systems, while not replacing human assessors (see Section 3.3), can complement human assessments.

Disclaimer/funding

This research was part of a larger study on the feasibility of introducing a large-scale system of essentiality assessments, funded by the European Commission's Joint Research Centre, JRC/SVQ/2018/B.6/0024/OC. The larger study is published as Bekkers et al. (2020a).

CRediT authorship contribution statement

Rudi Bekkers: Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing, Funding acquisition. **Elena M. Tur:** Methodology, Formal analysis, Investigation, Writing – original draft, Writing – review & editing. **Joachim Henkel:** Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing, Funding acquisition. **Tommy van der Vorst:** Conceptualization, Methodology, Investigation, Writing – original draft, Writing – review & editing, Funding acquisition. **Menno Driesse:** Investigation, Writing – original draft, Writing – review & editing. **Jorge L. Contreras:** Writing – original draft, Writing – review & editing.

Declaration of competing interest

All authors report financial support was provided by European Commission Joint Research Centre.

Rudi Bekkers reports a relationship with European Patent Office that includes: consulting or advisory. Rudi Bekkers reports a relationship with European Telecommunications Standards Institute that includes: consulting or advisory. Rudi Bekkers reports a relationship with IEEE that includes: travel reimbursement. Rudi Bekkers reports a relationship with Apple Inc. that includes: paid expert testimony. Rudi Bekkers reports a relationship with HTC America Inc. that includes: paid expert testimony. Rudi Bekkers reports a relationship with TCL Corp that includes: paid expert testimony.

Jorge Contreras reports a relationship with IETF that includes: consulting or advisory. Jorge Contreras reports a relationship with IEEE that includes: travel reimbursement. Jorge Contreras reports a relationship with ANSI that includes: board membership. Jorge Contreras reports a relationship with IRS that includes: paid expert testimony. Jorge Contreras reports a relationship with Thales that includes: paid expert testimony. Jorge Contreras reports a relationship with Lenovo that includes: paid expert testimony. Jorge Contreras reports a relationship with Samsung that includes: paid expert testimony.

Joachim Henkel reports relationships with Apple and various litigants in SEP-related court cases that includes: paid expert testimony.

Acknowledgements

We would like to thank the six European patent offices that made their examiners available to carry out essentiality assessments in our pilot experiment. We are also very grateful to the SEP owners and others involved in this project through discussions and providing the necessary information for the experiment. We specifically thank Christian Loyau, Edmund Mangold, Yann Ménière, Ruud Peters, Serge Raes, and Mats Sågfors for their valuable advice, patent pool administrators Avanci, One-Blue, SISVEL, Via Licensing and 3G3P for sharing their extensive experience with essentiality assessments, and the European Commission for funding and guidance.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.respol.2022.104590.

References

Aoki, R., Arai, Y., 2018. Strategic Declaration of Standard Essential Patents. In: RIETI Discussion Paper Series 18-E-035, Retrieved 19/02/22. RIETI. https://www.rieti.go.jp/jp/publications/dp/18e035.pdf.

Aoki, R., Nagaoka, S., 2005. Coalition formation for a consortium standard. Economic Review (Keizai Kenkyu) 55 (4), 345–356.

Arora, A., Ceccagnoli, M., 2006. Patent protection, complementary assets, and firms' incentives for technology licensing. Manag. Sci. 52 (2), 293–308. https://doi.org/10.1287/mnsc.1050.0437.

Arora, A., Fosfuri, A., Gambardella, A., 2001a. Markets for technology and their implications for corporate strategy. Ind. Corp. Chang. 10 (2), 419–451. https://doi.org/10.1093/icc/10.2.419.

Arora, A., Fosfuri, A., Gambardella, A., 2001b. Markets for Technology: The Economics of Innovation and Corporate Strategy. MIT Press, Cambridge, MA.

Arts, S., Cassiman, B., Gomez, J., 2017. Text matching to measure patent similarity. Strateg. Manag. J. 39 (1), 62–84. https://doi.org/10.1002/smj.2699.

Baron, J., Pohlmann, T., 2018. Mapping standards to patents using declarations of standard-essential patents. J. Econ. Manag. Strateg. 27 (3), 504–534. https://doi.org/10.1111/jems.12255.

Baron, J., Contreras, J.L., Husovec, M., Larouche, P., 2019. In: Thumm, N. (Ed.), Making the Rules: The Governance of Standard Development Organizations and their Policies on Intellectual Property Rights. Publications Office of the European Union, Luxembourg. https://doi.org/10.2760/207268.

Beggs, A., 1992. The licensing of patents under asymmetric information. Int. J. Ind. Organ. 10 (2), 171–191. https://doi.org/10.1016/0167-7187(92)90014-P.

Bekkers, R., Henkel, J., Tur, E.M., van der Vorst, T., Driesse, M., Kang, B., Martinelli, A., Maas, W., Nijhof, B., Raiteri, E., Teubner, L., 2020a. Pilot study for essentiality assessment of Standard Essential Patents. In: Thumm, N. (Ed.), EUR 30111 EN / JRC119894. Publications Office of the European Union, Luxembourg, ISBN 978-92-76-16666-5. https://doi.org/10.2760/68906.

Bekkers, R., Raiteri, E., Martinelli, A., Tur, E., 2020b. Landscape Study of Potentially Essential Patents Disclosed to ETSI. Thumm, N. (ed.), JRC121411, Publications Office of the European Union, Luxembourg, doi:10.2760/313626, ISBN 978-92-76-20786-3".

Bekkers, R., Updegrove, A., 2013. A Study of IPR Policies and Practices of a Representative Group of Standards Setting Organizations Worldwide. Retrieved 19/02/22. National Academies of Science, Washington. https://www.nap.edu/resource/18510/Bekkers-Updegrove%20Paper_092013.pdf.

Bekkers, R., West, J., 2009. The limits to IPR standardization policies as evidenced by strategic patenting in UMTS. Telecommun. Policy 33 (1–2), 80–97. https://doi.org/10.1016/j.telpol.20 08.11.0 03.

Bessen, J., Meurer, M.J., 2009. Patent Failure: How Judges, Bureaucrats, and Lawyers Put Innovators at Risk. Princeton University Press.

Biddle, B., White, A., Woods, S.. How Many Standards in a Laptop? (And Other Empirical Questions). Retrieved 1/07/22. https://ssrn.com/abstract=1619440.

Brachtendorf, L., Gaessler, F., Harhoff, D., 2020. Truly Standard-Essential Patents? A Semantics-Based Analysis. CEPR Discussion Paper No. DP14726. Retrieved 19/02/22. https://ssrn.com/abstract=3603956.

Caves, R., Crookell, H., Killing, J., 1983. The imperfect market for technology licenses. Oxf. Bull. Econ. Stat. 45 (3), 249–267. https://doi.org/10.1111/j.1468-0084.1983.mp45003002.x.

Charles River Associates (CRA), 2016. Transparency, Predictability, and Efficiency of SSO-Based Standardization and SEP Licensing. A Report for the European Commission. Retrieved 19/02/22. https://ec.europa.eu/docsroom/documents/20506/attachments/1/translations/en/renditions/native.

Contreras, J.L. (Ed.), 2007. Standards Development Patent Policy Manual. American Bar Association (ABA), Committee on Technical Standardization, Section of Science & Technology Law. ABA Publishing, Chicago.

Contreras, J.L., 2013. Fixing FRAND: a pseudo-pool approach to standards-based patent licensing. Antitrust Law J. 79 (1), 47–97.

Contreras, J.L., 2017. Essentiality and standards-essential patents. In: Contreras, J. (Ed.), The Cambridge Handbook of Technical Standardization Law: Competition, Antitrust, and Patents. Cambridge University Press, New York, pp. 209–230. https://doi.org/10.1017/9781316416723.

Contreras, J.L., 2019. Technical standards, standards-setting organizations and intellectual property: A survey of the literature (with an emphasis on empirical approaches). In: Depoorter, B., Menell, P., Schwartz, D. (Eds.), Research Handbook on the Economics of Intellectual Property Law, Vol 2 – Analytical Methods. Edward Elgar, Cheltenham and Northampton, pp. 185–235. https://doi.org/10.4337/9781789903997.

Cyber Creative Institute, 2011. Evaluation of LTE Essential Patents Declared to ETSI. Retrieved 19/02/22. http://www.cybersoken.com/file/lte01EN.pdf.

Den Uijl, S., Bekkers, R., de Vries, H., 2013. Managing intellectual property using patent pools: lessons from three generations of pools in the optical disc industry. Calif. Manag. Rev. 55 (4), 31–50. https://doi.org/10.1525/cmr.2013.55.4.31.

Dewatripont, M., Legros, P., 2013. "Essential" patents, FRAND royalties and technological standards. J. Ind. Econ. 61 (4), 913–937. https://doi.org/10.1111/joie.12033.

European Commission, 2017. Communication from the Commission to the European Parliament, the Council and the European Economic and Social Committee: Setting out the EU approach to Standard Essential Patents. COM(2017) 712. Retrieved 19/02/22, https://ec.europa.eu/docsroom/documents/26583.

European Commission, 2020a. Communication from the Commission to the European Parliament, the Council, the European Economic and Social Committee and the Committee of the Regions: Making the most of the EU's innovative potential: An intellectual property action plan to support the EU's recovery and resilience. Brussels, 25.11.2020 COM(2020) 760 final. Retrieved 19/02/22, https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52020DC0760.

European Commission, 2020b. Press release: Speech by Executive Vice-President Vestager at the Press Conference on the Data Governance Act and the Action Plan on Intellectual Property. Brussels, 25 November 2020. Retrieved 19/02/22. European Commission, Brussels. https://ec.europa.eu/commission/presscorner/detail/en/speech_20_2210.

European Telecommunications Standards Institute (ETSI), 2021. Rules of Procedure of the European Telecommunications Standards Institute, Version Approved by a

Specially Convened Meeting of the ETSI General Assembly (1 December 2021). Retrieved 19/02/22. ETSI: Sofia Antipolis. https://www.etsi.org/images/files/IPR/etsi-ipr-policy.pdf.
Fairfield Resources International, 2007. Analysis of Patents Declared as Essential to GSM as of June 6, 2007. Retrieved 30/01/19. http://www.frlicense.com/recent.html.
Fairfield Resources International, 2009. Review of Patents Declared as Essential to WCDMA Through December, 2008. Retrieved 30/01/19. http://www.frlicense.com/recent.html.
Fairfield Resources International, 2010. Review of Patents Declared as Essential to LTE and SAE (4G Wireless Standards) Through June 30, 2009. http://www.frlicense.com/recent.html.
Gallini, N., Wright, B., 1990. Technology transfer under asymmetric information. RAND J. Econ. 21 (1), 147–160. https://doi.org/10.2307/2555500.
Gans, J., Hsu, D., Stern, S., 2008. The impact of uncertain intellectual property rights on the market for ideas: evidence from patent grant delays. Manag. Sci. 54 (5), 982–997. https://doi.org/10.1287/mnsc.1070.0814.
Gilbert, R., 2017. Collective rights organizations: A guide to benefits, costs and antitrust safeguards. In: Contreras, J. (Ed.), The Cambridge Handbook of Technical Standardization Law: Competition, Antitrust, and Patents. Cambridge University Press, p. 125. https://doi.org/10.1017/9781316416723.011.
Goldstein, L., Kearsey, B., 2004. Technology Patent Licensing. Aspatore Books.
Goodman, D., Myers, R., 2005. 3GPP Cellular Standards and Patents. IEEE WirelessCom 2005.
Henkel, J., 2022. Licensing standard-essential patents in the IoT – A value chain perspective on the markets for technology. Res. Policy in print.
Henkel, J., Zischka, H., 2019. How many patents are truly valid? Extent, causes, and remedies for latent patent invalidity. Eur. J. Law Econ. 48 (2), 195–239. https://doi.org/10.1007/s10657-019-09627-4.
IPLytics, 2021. AI May be the Solution to Skyrocketing Numbers of SEP Declarations. Retrieved 19/02/22. https://www.iplytics.com/report/ai-solution-skyrocketing-numbers-sep-declarations/.
iRunway, 2012. Patent & Landscape Analysis of 4G-LTE Technology. Retrieved 30/01/19, https://www.i-runway.com/images/pdf/iRunway - Patent & Landscape Analysis of 4G-LTE.pdf.
Jaffe, A., Lerner, J., 2006. Innovation and its discontents. NBER Innov. Policy Econ 6 (1), 27–65.
Jefferies & Company, 2011. Research in Motion (RIMM) Limited patent value; cut target to salvage value of $21 as we wait for QNX. Retrieved 19/02/22. https://ipcloseup.files.wordpress.com/2011/10/rimm.pdf.
Lamoreaux, N., Sokoloff, K., 1999. Inventive Activity and the Market for Technology in the United States, 1840–1920. NBER Working Paper 7107. National Bureau of Economics Research, Inc, Cambridge, MA.
Layne-Farrar, A., Lerner, J., 2011. To join or not to join: Examining patent pool participation and rent sharing rules. Int. J. Ind. Organ. 29 (2), 294–303. https://doi.org/10.1016/j.ijindorg.2010.08.006.
Lemley, M., Shapiro, C., 2007. Patent holdup and royalty stacking. Texas Law Review 85, 1991–2048.
Lemley, M., Simcoe, T., 2019. How essential are standard-essential patents? Cornell Law Review 104 (3), 607–642.
Macho-Stadler, I., Martinez-Giralt, X., Perez-Castrillo, J., 1996. The role of information in licensing contract design. Res. Policy 25 (1), 43–57. https://doi.org/10.1016/0048-7333(94)00815-9.
Madiès, T., Guellec, D., Prager, J. (Eds.), 2014. Patent Markets in the Global Knowledge Economy: Theory. Empirics and Public Policy Implications, Cambridge University Press.
Martin, D., De Meyer, C., 2006. Patent Counting, a Misleading Index of Patent Value: A Critique of Goodman & Myers and its Uses. Retrieved 19/02/22. https://ssrn.com/abstract=949439.
Merges, R., Mattioli, M., 2017. Measuring the costs and benefits of patent pools. Ohio State Law J. 78 (2), 282–347.
Milgrom, P., Roberts, J., 1992. Economics, Organization and Management. Prentice-Hall, Englewood Cliffs, NJ.
Natterer, M., 2016. Ähnlichkeit von Patenten: Entwicklung, empirische Validierung und ökonomische Anwendung eines textbasierten Ähnlichkeitsmaßes. Retrieved 19/02/22. Verlag für Nationalökonomie, Management und Politikberatung.
PA Consulting Group, 2015. LTE Essential IPR. PA's 3GPP-LTE Database and Report.
Partners, Article One, 2011. LTE Standard Essential Patents Now and in the Future. Retrieved 30/01/19. http://newsletters.articleonepartners.com/news_f1317eac-ee13-5a66-d0f5-38ea99a4c1eeLTE-Standard-Essential-Patents-Now-and-in-the-Future.pdf.
Qualcomm, 2006. Proposal for Defining Articulated Actual Problems and Solutions (DAAPS), Document Submitted to ETSI as Document ETSI GA/IPRR04(06)13, 10 May 2006.
Scherer, F.M., Harhoff, D., 2000. Technology policy for a world of skew-distributed outcomes. Res. Policy 29 (4–5), 559–566. https://doi.org/10.1016/S0048-7333(99)00089-X.
Siebrasse, N., Cotter, T., 2017. Judicially determined FRAND royalties. In: Contreras, J. (Ed.), The Cambridge Handbook of Technical Standardization Law: Competition, Antitrust, and Patents. Cambridge University Press, pp. 365–388. https://doi.org/10.1017/9781316416723.
Younge, K., Kuhn, J., 2016. Patent-to-Patent Similarity: A Vector Space Model. Retrieved 19/02/22. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2709238.