



Contents lists available at ScienceDirect

Government Information Quarterly

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

Decentralising the patent system

Gaétan de Rassenfosse^a, Kyle Higham^{b,*}^a Ecole polytechnique fédérale de Lausanne, Switzerland^b Institute of Innovation Research, Hitotsubashi University, Japan

A B S T R A C T

Modern patent systems are slow, inefficient, expensive, and may result in outcomes that actively harm technological progress. This paper proposes a substantive rethink of these systems and lays a foundation upon which practical solutions can be built. Many solutions proposed in the past, such as prior-art bounties, outsourced examination, and dynamic fee setting, have gone unheeded due to the cost of administering them and the rigidity of the patent system. We explore how distributed ledger technologies (DLTs) enable these major changes by altering the way stakeholders are able to interact with the patent records system. We find that transitioning to a DLT-based patent records system can enable many previously suggested improvements to current patent systems in a flexible, scalable, and transparent manner. The case for such a transition is strengthened when jointly considering the complex but common roots of problems facing modern patent systems, rather than a balkanised set of technical solutions to address each issue independently. Noting that a DLT-based system is not a panacea, we also provide comment on the political, legal, and organisational challenges that must be overcome for such changes to be implemented at scale.

1. Introduction

The core rationale behind the patent system has not changed since 15th century when the first well-documented patent system was concocted in Venice (Long, 1991; Prager, 1944; Sichelman & O'Connor, 2012). By granting inventors the right to exclude others from a newly discovered slice of technological space, the patent system provides an incentive to pour more resources into the inventive process. Furthermore, the disclosure of the technical details of the invention in the patent document is supposed to facilitate knowledge diffusion. However, it has been argued that current patent systems are not a particularly efficient means of achieving these goals (Bessen & Meurer, 2008; Jaffe & Lerner, 2004; Stiglitz, 2007; Thurow, 1997), and additionally come with such high costs to society that many have wondered why governments should grant these rights at all (Boldrin & Levine, 2008, 2013; Mirowski, 2011).

The shortcomings of modern patent systems are well-documented. We propose a reconfiguration of patent record systems that facilitates solutions to these problems through decentralisation of patent-related processes and information, in conjunction with increased transparency, flexibility, and efficiency. In particular, we ask whether distributed ledger technologies (DLTs) constitute an appropriate technical foundation for such a change, and consider how they may be

practically implemented. The potential for DLTs to play this role is illustrated through a series of proposals. These proposals, many of which are already discussed in the literature, are aimed at alleviating well-known problems associated with patent systems, including long grant lags (Ackerman, 2011; Gans, Hsu, & Stern, 2008), high levels of patent invalidity (Frakes & Wasserman, 2017; Lemley & Shapiro, 2005), anti-competitive patenting strategies (Blind, Cremers, & Mueller, 2009; Macdonald, 2004), the high risk and cost of litigation (Bessen, Ford, & Meurer, 2011; Bessen & Meurer, 2013; Lemley, Richardson, & Oliver, 2018), and many other problems that have been covered in recent years (Eckert & Langinier, 2014; Hall & Harhoff, 2012; Sampat, 2018). We explore many of these issues throughout this work, and explain how a modernized patent system can address them.

It has become clear in recent years that it is not merely inefficiencies within patent offices that lie at the root of these problems, but the political and legal systems to which these offices are duty-bound. It is easy to forget that patent offices are resource-constrained actors with little control over the legal and industrial environments in which they operate. Financial constraints and extended patent pendency times, the decisions around patent-eligible subject matter, or even differences in novelty and non-obviousness thresholds between the patent office and the courts, are crucial issues with current patent systems that are often entirely out of the control of patent offices (Jaffe & Lerner, 2004; Rai, 2008; Sohn, 2011).^{2,3}

* Corresponding author.

E-mail addresses: gaetan.derassenfosse@epfl.ch (G. de Rassenfosse), higham@iir.hit-u.ac.jp (K. Higham).¹ Most of this research was conducted while this author was working at École polytechnique fédérale de Lausanne.² Of course, many of the benefits and drawbacks of current patent systems lie outside of these bureaucracies, and any proposed changes to patent systems must also consider the flow-on effects that these policies may have.³ In the United States, however, there have been recent developments towards increasing patent office autonomy (Wasserman, 2012).<https://doi.org/10.1016/j.giq.2020.101559>

Received 6 July 2020; Received in revised form 18 December 2020; Accepted 20 December 2020

0740-624X/© 2020 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

By decentralising examination processes, facilitating information sharing, and increasing transparency of these operations, we argue that, despite these restrictions, patent offices will be more likely to grant valid patents in a timely fashion and simultaneously mitigate the negative externalities associated with current patent systems.

We use the word ‘decentralisation’ very broadly in this work. First, *patent data* and the management of these data are currently quite centralised; patent offices hold and organise all patent information, and we generally trust them with this responsibility. However, this trust may not hold across all jurisdictions or, indeed, indefinitely. Centralisation also appears to be simply unnecessary; most patent information is public by nature. A more decentralised source appears to be sensible for this purpose. Second, and perhaps most importantly, the *patenting process* is even more centralised than the data. Patent offices represent a bottleneck in innovation systems more widely, and this situation can be largely explained by the centralisation of this system and its interaction with public policy (Jaffe & Lerner, 2004). The primary function of patent offices today is to examine patent applications; however, this process is extremely slow and often results in poor-quality output. At the same time, there exists a large pool of expertise outside local patent offices that is not being leveraged to address these problems—decentralising some parts of the patenting process to allow the patent office to tap into this collective intelligence may be sensible (Nam, 2012; Noveck, 2006).

It is also necessary to note the distinction between decentralisation of processes and the distribution of information across a network. In this work, we advocate for both a decentralisation of *processes* as described above and a decentralisation of *responsibility* for a shared records system. We foresee an international DLT-based system that may be utilised by many independent jurisdictions in accordance with their local patent laws, policies, and procedures. For this reason, and to ensure a high level of transparency, power would need to be distributed across multiple jurisdictions for this vision to be feasible. *Distribution* of computing power and data storage are clearly important considerations for such a system, but these issues are not central to our arguments for decentralisation as described above. For this reason, while DLTs are ‘distributed’ in name and nature, these distributive properties are thought of as tools that facilitate decentralisation.

This article assumes that the reader understands the fundamental mechanics and properties of distributed ledgers. Appendix A provides a brief overview of these concepts. Readers can find a detailed description of this class of technologies in the review by Rauchs et al. (2018), as well as comprehensive investigations on use-cases and wider economic implications of these technologies in the books by Swan (2015) and Berg, Davidson, and Potts (2019). However, it is important to keep in mind that these technologies are still new, and different configurations continue to be invented and implemented for new use-cases. For this reason, we do not wish to restrict the proposals detailed in this work to one particular system architecture, and thus will continue to use the general term DLT. Blockchain, for example, was the first and is the most widely used form of DLT, but there exists an entire ecosystem of DLTs that differ from each other in many fundamental ways. Some of these blockchain-alternatives aim to fix particular shortcomings of blockchain to make them more suited to particular use-cases. Usually these improvements are made by facilitating both vastly increased transaction rates on the network as well as scalability of the DLT (El Ioini & Pahl, 2018; Pervez, Muneeb, Irfan, & Haq, 2018; Yang, Garg, Raza, Herbert, & Blockchain, 2018).⁴ Further, the most appropriate system for any particular use-case, including one designed with patent systems in mind, may not yet exist.

⁴ For example, directed acyclic graph-based DLTs (Benčić & Žarko, 2018) are sometimes touted as the natural successors of blockchains (Pervez et al., 2018; Popov, 2016), and are already gaining traction as an appropriate form of DLT for applications such as microtransactions (Popov, 2016).

The structure of this document is as follows. We first explore in Section 2 the roots of the most significant problems plaguing modern patent systems, the necessary characteristics for technical solutions to these problems, and the degree to which DLTs match these characteristics. The main sections of the paper, Sections 3 and 4, split potential DLT-enabled changes to patent systems into those that run complementary to current procedures and those that fundamentally restructure them. Both of these sections are organised chronologically through the life of an invention. Section 5 outlines the important new roles and responsibilities that patent offices would need to assume to maintain and regulate functional patent systems of the future, should they become more decentralised and dynamic. Section 6 discusses the legal, technical, and regulatory challenges that need to be overcome before a DLT-based system could be implemented at scale. Section 7 concludes.

2. Addressing problems at their roots

The roots of many of the major problems with current patent systems are shared and intermingled. As such, we will start by characterising these roots, some of which may appear relatively benign, to illuminate the primary targets at which any solution should take aim. First, patent offices are very slow to grant patents. The specific reasons for this slowness differ by jurisdiction (Mejer & van Pottelsberghe, 2011) but generally lie at the intersection of firm strategy, patent policy, and patent office resourcing. This lengthy delay is associated with numerous negative impacts that accompany uncertain intellectual property (IP) rights (Czarnitzki, Hussinger, & Schneider, 2015; Gans et al., 2008; Heger & Hussinger, 2017). Second, patent offices grant many invalid patents. While likelihood of invalidity varies by office (de Rassenfosse, Griffiths, Jaffe, & Webster, 2020), patent examiners have a limited amount of time to search for prior art that may impact the patentability of any individual application, and there is an enormous amount of public information to sift through. Prior art that would result in an emphatic rejection of a patent application is inevitably missed, and the resulting ‘weak’ patents may block follow-on innovation, cause unanticipated financial harm to the patent owner, and generally lead to distrust in patent systems for protecting IP. Third, although a significant amount of information and workload sharing between jurisdictions does occur, particularly for applications that pass through the Patent Cooperation Treaty route, a global patent records system is still a long way away. The ‘patent prosecution highway’ is a step in the right direction, but remains a complex bundle of bilateral agreements to share examination workloads (Pitts & Kim, 2009). Prior art is *not* jurisdiction specific, so this aspect of patent examination has particularly great potential for efficiency improvements. Lastly, patent prosecution processes are largely disconnected from the inventive process that precedes them and the commercialisation processes that follow them, both of which are extremely trust-dependent. Third parties have started to offer DLT-based solutions to manage these latter external processes in recent years.⁵ The World Intellectual Property Organisation (WIPO) has recently started offering a similar service (WIPO PROOF)⁶ which, while not DLT-based, could easily be incorporated into a DLT records system. Crucial in all of these processes is IP ownership, which is highly opaque and presents a significant trust barrier with regard to markets for technology (Jensen, Palangkaraya, & Webster, 2015).

2.1. Characterising a solution

The problems described above share some intermingled roots that are amenable to technical solutions, namely: centralised examination

⁵ See, e.g., bernstein.io and ipwe.com, respectively. (Last accessed 10/12/2020.)

⁶ wipo.int/wipoproof/en/index.html (Last accessed 10/12/2020.)

as a hindrance, inadequate prior art search, barriers to the sharing of workloads and information between jurisdictions, and various aspects of trust throughout the life of an invention (pre-application, during prosecution, and post-grant). We can now start to think about frameworks that may offer a foundation on which solutions to these problems can be built. Many of these problems exist in patent systems all over the world, and would be mitigated by higher levels and new modes of cooperation. As such, a new system must allow cooperation (such as information and workload sharing) between jurisdictions while also remaining flexible enough for participating offices to retain autonomy and experiment with solutions to their (potentially unique) problems (Ølne, Ubacht, & Janssen, 2017). We suggest that current, mostly centralised, approaches to patent examination (that is, the idea that all components of patent prosecution must be conducted in-house) are unnecessarily limiting. Opening up aspects of this process to third parties (including both foreign patent offices and local experts), from prior art search to substantive examination, requires a system that is transparent by design and reduces reliance on trust between parties involved. The same is true for modern markets for technology. Solutions relying on traditional technical systems may currently demand an unacceptable level of trust and are less amenable to extending accessibility for implementation on a truly global scale. Of course, applicants usually trust patent offices with their valuable IP, but this is not the only interface within patent systems where trust is a concern.

For the purposes of this article, we consider two distinct forms of trust based on whether or not malice is at play. In the first instance, for example, it is the nature of patents to be public knowledge, and it is the nature of knowledge to be non-rival, so identity and ownership must play a central role in a trustworthy market for technology—this environment is perfect for bad actors with the intent to deceive. More perniciously, recent work has uncovered significant bias against foreign patent applicants in technological fields considered ‘strategic’ for the jurisdiction of the receiving office (de Rassenfosse & Raiteri, 2020). In the second instance, potential patentees pay patent offices a significant amount of money to obtain and maintain a patent—they *trust* that the patent office is competent and that their patent is valid. If this is not the case, it was not an intentional act of deceit on the part of the patent office, it was merely a mistake. Further, patent offices do not necessarily trust the completeness of prior art search conducted by other offices, which are often used as reference in applications that are made to multiple jurisdictions. To increase completeness of, and confidence in, search reports, the IP5 launched the Collaborative Search and Examination pilot program in 2018.^{7,8} This program, requiring independent prior art search and coordination across offices is of precisely the kind that could be integrated into a DLT-based patent information system, with potential for relatively straightforward scaling to more offices. A specific system for doing so has even been recently proposed (Marschall, 2020). In summary, a system that introduces a significant level of transparency, while facilitating the cooperation of entities with common goals and aligned incentives, not only mitigates acts of intentional deceit, but may also reveal mistakes (including errors of omission) early enough to rectify them.

2.2. Do we need a DLT?

DLTs comprise a relatively new class of technology and, as with any potential replacement to tried-and-tested traditional systems, it is important to think carefully about the necessity of such a transition. The

first question usually associated with this concern is: Do we need a blockchain? (Wüst & Gervais, 2018). Of course, we do not wish to restrict ourselves to the technical framework of blockchain, but the sentiment holds: Do we need a DLT?

It would be absurd to claim that *every* problem present in modern patent systems requires a DLT-based solution; many can, in fact, be solved by the adoption of more traditional technological tools. We do not claim that DLTs are a panacea. However, in light of the desired properties outlined above that we would expect to find in a solution, we suggest that DLTs are a suitable class of candidate systems. Not only do DLTs provide a level of transparency and trust that are not inherent to alternatives, but they are a highly flexible enabling technology into which novel features (which may take advantage of DLT-based tools such as so-called ‘smart contracts’) can be introduced as IP policy evolves. Implementing such a system will also mitigate the balkanisation of the various technical systems and processes used in patent offices around the world (such as the aforementioned WIPO PROOF) and may even act as an impetus for their international harmonisation. We argue that if a few important problems are best solved with a DLT-based system (and we believe this to be the case), then we should also explore how other problems may be solved within the same framework, even if they can also be solved with independent tools (each with their own maintenance costs) at the detriment to the coordination of the system as a whole. While the most comprehensive proposals in this work focus on DLT-based solutions to the examination bottleneck and incomplete prior art search, a number of auxiliary applications are briefly explored in order to demonstrate the *adaptability* of a DLT to handle such processes in the context of patent systems (Ølne & Jansen, 2017). We additionally assert that DLTs not only streamline many processes within the patent office as it is organised at present, but they also significantly expand the realm of possibilities for improvements to current policies and protocols. Many of these possibilities may indeed involve interactions between parties that do not necessarily trust each other (in either sense) such as contracting out patent examination (see Section 4.2).⁹

We tentatively suggest that such a DLT-based system should be of the public-permissioned variety. In our suggested configuration wherein nodes are patent offices around the world, the nodes have aligned incentives to secure the network and keep a reliable record, including both the utility they obtain from its existence as well as their own reputation. This means that there is no need to incentivise those outside the patent offices (via, for example, cryptocurrencies or tokens) to secure the network, and, in fact, no reason for the DLT to be permissionless at all, as long as clients can interact with it and the information on the DLT itself is public (albeit encrypted when necessary). Further, a permissioned blockchain governed by a consortium of offices allows much more flexibility in the design and use of the system—particularly with regard to the autonomy of individual offices (Kannengießer, Lins, Dehling, & Sunyaev, 2020). For a DLT-based patent records system to be politically palatable, patent offices as a collective need to be able to have control over its use and permissioned DLTs are more amenable to such an arrangement (Zavolokina, Ziolkowski, Bauer, & Schwabe, 2020), both at its genesis and in the future. Permissioned DLTs also avoid some of the downsides often associated with permissionless systems, such as transaction throughput and security. In particular, transaction frequency is no longer a bottleneck, as consensus protocols on permissioned DLTs are very efficient relative to current permissionless protocols (Sousa, Besani, & Vukolic, 2018); and validator nodes are known, mitigating the risk of attacks that aim to control consensus. The security of any private information (that is, information not recorded in non-encrypted form on

⁷ The IP5 is a group of the five largest patent offices in the world (United States, Europe, Japan, China, and South Korea) which cooperate to reduce duplicative work related to patent prosecution, among other goals such as harmonisation of patent policy and improved data accessibility.

⁸ <https://www.wipo.int/pct/en/filing/cse.html> (Last accessed 10/12/2020)

⁹ It is worth noting that trust is not the only factor driving the adoption of DLTs; after all, firms currently use DLT-based records internally. The benefits of DLTs for streamlining asset management (Savelyev, 2018), data sharing (Xia et al., 2017), and identity management (Dunphy & Petitcolas, 2018), for example, are well-known.

the ledger) is equivalent to that of traditional database systems.

The adoption of a system that allows for a reliable decentralisation of patent prosecution processes will be associated with lower per-patent examination costs for patent offices or lower societal cost due to reduced patent invalidation rates, or both. In addition, the flexibility to experiment with and augment programs designed to increase patent office income or savings will in all likelihood outweigh both upfront costs and maintenance costs in the long-run. It is vital to keep in mind that patent offices are public institutions, so benefits to society as a whole must be included in any cost-benefit analysis—the majority of our proposals here are put forward with societal benefit as a priority. By leveraging the advantages of DLTs over traditional means of record-keeping, we can re-frame the primary functions of government-run patent systems (Ølnes et al., 2017) to ensure that these benefits are realised.

3. DLTs for unlocking flexibility

In the minds of many, DLTs have the potential to change the nature of institutional trust and offer the opportunity to replace traditional institutions with decentralised systems entirely. The abolition of patent offices, or even patents themselves, is not the topic of the current work. In fact, even in a much more decentralised IP system, we see patent offices as essential players. By nature, a patent bestows upon the owner a government-granted right—as long as patents exist, the government has the responsibility to ensure they work as intended. This regulatory role will remain paramount, and we discuss this role in more detail in Section 5.

In this work, we make a distinction between those proposals that leverage DLT-based records for structural changes to the patenting process and those that may raise the efficiency of current procedures. In this section, we discuss the latter, the less dramatic changes that patent offices can make that may nonetheless have significant benefits to the functioning of the patent system as a whole. The following proposals may run *complementary to current office procedures* and facilitate greater experimentation with new policies than is currently possible. Here, we think of DLTs as tools for ‘unlocking flexibility,’ both for patent offices (by facilitating policy experimentalism (Gong, Yang, & Shi, 2020)) and users (through new technology management tools).

3.1. Patent application

The first, and perhaps most obvious, step towards a DLT-based patent records system is to register patent applications on a DLT. This process is a small extension to current procedures: it merely adds the record to the ledger to provide notice that an application has been submitted to the patent office. Additionally, all future transactions related to this application and, if successful, subsequent patent can be recorded on the DLT as well, providing full transparency and immutability. At the United States Patent and Trademark Office (USPTO), this record-keeping already takes place through PAIR,¹⁰ while the European Patent Office (EPO) achieves much the same level of transparency through the European Patent Register.¹¹ The *contents* of the information placed on the DLT in this manner does not have to be public, however. For example, all information regarding a new patent application (and any amendments to this application) can be hashed before being recorded on the DLT,¹² and decrypted (as a new record) when it is made public as a typical pre-grant application.

Functionally, this procedure is identical to current procedures. At the

time of decryption, the transaction history related to the application can also be made public. This process ensures that the entire history of the application or patent becomes eventually public information and, if recorded correctly, can be trusted to be accurate due to the nature of the DLT.¹³ While we see no reason to distrust the patent office to perform these tasks, the present mechanism allows technology transfer between potentially untrusting parties while the application is still secret—as the hash allows for easy verification that the application being licensed pre-grant or otherwise transferred is the same as the application that is awaiting examination (without requiring any affirmation from the patent office).

3.2. Patents as assets

Patent rights are part of a bargain, whereby an applicant discloses the details of an invention, and the government grants the applicant temporary rights to exclude its use by others. After this point in time, the applicant may treat this IP as any other asset—they can license its use to others, sell it outright, securitise or collateralise it, or even place it directly into the public domain.

As such, perhaps the most convincing non-structural rationale for transitioning to a DLT-based patent records system is how ownership and use of these assets may be recorded and managed by both patent offices and their users. DLTs also provide many opportunities to third-parties that offer auxiliary services to both potential applicants and patent owners. These opportunities are numerous and cover many aspects of the innovation process, including trade secrets, collaboration, and supply chain management. However, here, we will focus on the direct and immediate benefits of DLT-based patent records: notice of patent ownership, ease of licensing, marking, and the automatic public disclosure of these actions. Third parties have already begun to offer private opt-in versions of some of these services; indeed, they have emerged in the absence of any patent office initiatives to incorporate DLTs into their processes.¹⁴

3.2.1. Patent ownership

In most jurisdictions, there is no legal obligation to provide the patent office with information about changes in patent ownership. This ‘notice failure’ introduces significant, and many would argue unnecessary, inefficiencies into the innovation system (Menell & Meurer, 2013). In many cases, even the government itself does not know exactly to whom these government-granted rights belong. This problem is well-known—one effort from the USPTO to attempt to remedy this situation was eventually abandoned due to concerns about undue administrative burden on patentees (Anderson, 2015).¹⁵ DLT-based records are inherently compatible with a requirement that any change of ownership (including to applications) needs to be registered with the patent office within a reasonable time frame from the transfer date. All parties involved need simply to provide digital signatures to authorise the change, without having to involve the patent office directly.¹⁶ This metadata would automatically be updated on the patent record. Many firms *currently* report reassignments to the patent office (Graham, Marco, & Myers, 2018)—DLT-based reassignment could streamline this process.

The USPTO’s proposed change mentioned above also required that all entities that may be able to *enforce* a particular patent be listed on the

¹⁰ <https://portal.uspto.gov/pair/PublicPair> (last accessed 10/12/2020)

¹¹ <https://register.epo.org/regviewer> (last accessed 10/12/2020)

¹² Cryptographic hashing is a process that takes arbitrary data as input and that outputs an apparently random string of ones and zeros with a fixed length. Crucially, this function always gives the same output for a given input, and it is impossible to reverse engineer the original data from this output.

¹³ Knowing the method of encryption as well as the original text, it is straightforward to check that the decrypted document was indeed what was encrypted in the first place.

¹⁴ See e.g., ipwe.com (last accessed 10/12/2020).

¹⁵ This view was predominantly put forward by large corporate entities with complex ownership structures, e.g. see <https://www.govinfo.gov/app/details/FR-2014-01-24/2014-01195> (last accessed 10/12/2020).

¹⁶ Note, however, that this kind of action should be subject to additional security measures, see Section 6.3.

document, along with any changes in these entities over time.¹⁷ This requirement appears to be the component of the proposal that would place the most ‘undue’ burden on patent holders. However, these potential enforcers are precisely the entities about whom competitors and potential infringers are concerned. Therefore, information on enforcers would likely go the furthest in addressing the inefficiencies associated with notice failure. This suggestion remains an appropriate one as part of broader reform efforts, as this burden is currently placed on competitors and not those who are responsible for the complexity of ownership structures in the first place—the patent holders themselves—and who are in the best position to provide this information (Anderson, 2015). Like ownership information, there is little justification for obscuring potential enforcement entities (and any changes to them) from the patent office and the public.¹⁸

DLTs would integrate nicely into any policy that would seek to clarify patent ownership or enforcement status at any particular point in time (Allen, Berg, Davidson, & Potts, 2020). These changes would be recorded in the patent metadata without directly involving the patent office. The firm that applied for the patent is listed on the application, the firm that was granted the patent is listed on the patent itself, and any transfers of ownership would be recorded (and time-stamped).¹⁹ These events can easily be aggregated to view a full ownership history. Transfers would *not* be, in the traditional sense, ‘reported’ to any authority—they are transactions between private parties made official by an associated record on the DLT. Patent offices, like the public, are simply observers in this scenario. When ownership is transferred, there is no doubt about the parties involved or the timing of the transaction, and the event is made immediately public.

3.2.2. Licensing and markets for technology

Scholars have pointed out that the market for technology, be it a license or transfer, is an extremely information-poor setting with regard to pricing (Arora & Gambardella, 2010; Lemley & Myhrvold, 2007). Unlike markets for other goods, negotiations for IP rights are often closed-door ordeals with significant information asymmetries between parties (Buenstorf & Geissler, 2012; Gallini & Wright, 1990). Details regarding patent ownership (or even patent *application* ownership) described above will also significantly increase trust in the market for technology (Jensen et al., 2015). Having certainty about who you are buying from, and that the rights being purchased are legitimate, are essential precursors to a functional market for non-rival goods. If patent rights (and their ownership status) are recorded in the same place regardless of the jurisdiction in which they are valid, this will further thicken the market and in many cases provide important price information; this feature could prove very useful to multinational firms, in particular.

Firms would benefit significantly from a thicker and richer (information-wise) market for technology. Perhaps most obviously, having access to a thick market of buyers and sellers allows more efficient distribution of knowledge resources across parties—on one side, buyers can access new technologies at market price, and on the other side, suppliers can raise funds by allowing others access to their inventions. In the latter case, these funds can be funnelled into other activities such as

follow-on innovation or investment in tangible capital (Arora, Fosfuri, & Gambardella, 2001; de Rassenfosse, 2012; Sichelman & O’Connor, 2012). Information on the value of the usage rights for particular inventions is also invaluable to all players in this market. In particular, knowing the market value of inventions would have a drastic impact on the pricing of IP rights, likely lowering licensing and transfer costs for everyone, increasing the pace of follow-on innovation, and reducing the impact of commercialisation-inhibiting phenomena such as patent thickets (Hall, Helmers, & von Graevenitz, 2015).

When patents are all recorded on a DLT, smart contracts can reduce the complexity of some simple forms of licensing and lower the barriers to licensing otherwise dormant IP (Gambardella, Giuri, & Luzzi, 2007).²⁰ Both the USPTO and WIPO facilitate some form of post-grant patent licensing market,²¹ and pre-written contracts could modernise these services: the licensor sets simple contractual terms, places the licenses in some central marketplace,²² and anyone who would like to license the technology could simply buy the contract.²³ This kind of system would incur little or no administrative costs on either party once the contract is digitally signed. Whether or not a patent is licensable, as well as a list of current licenses,²⁴ would be recorded in the DLT and aggregated with the other metadata of the patent in question. This type of licensing would probably work best when bundles of patents are offered (such as patent pools) or in the context of standard-essential patents (Lerner & Tirole, 2015; Shapiro, 2000).

All of these licensing and transfer transactions recorded on the DLT would constitute valuable and easily aggregated information to assist in the correction of the information asymmetry that exists within the current market.

3.3. Virtual patent marking

Virtual Patent Marks (VPMs) were introduced in the United States with the America Invents Act in 2012, which allows patent owners to indicate patent-product relationships on the internet as a complement to physically marking the product (de Rassenfosse, 2018; de Rassenfosse & Higham, 2020; USPTO, 2014). VPMs were introduced in the United Kingdom in 2014 under the name ‘webmarking.’ Marking serves several functions including, importantly in the current context, informing would-be imitators of patent rights that may be infringed should the product be copied. If the patent is marked and subsequently infringed, the patent owner can claim damages from the time the VPMs were posted online (as opposed to from the moment a formal infringement notice has been mailed to the alleged infringer). We suggest that virtual marking is perhaps under-utilised as a tool for not only monitoring commercialisation of patented inventions (which may in itself have important policy implications) but also as a way to incentivise proper disclosure.

A patentee could virtually mark their patents by linking them to their associated products via a transaction on the DLT.²⁵ This process would create a two-way link between patent and product: one could not only find the patents marked on the product page as is currently the case, but

¹⁷ See <https://www.govinfo.gov/app/details/FR-2014-01-24/2014-01195> for more information (last accessed 10/12/2020).

¹⁸ Note in the case of ownership, however, there is a significant incentive on the part of the new assignee to ensure the transfer is promptly reported to the patent office. If it is not, the transfer may be voided in favour of a subsequent transfer that does get reported. (See 35 USC §261).

¹⁹ Loss of ownership is not a concern in the same way that it is for cryptocurrencies (e.g. money.cnn.com/2013/11/29/news/bitcoin-haul-landfill) because ownership is not secret. A registration process for assignees means that there is an authority with whom to interface in case of security issues. For a brief overview of security issues in the context of DLT-based patent records, see Section 6.3.

²⁰ Patent licensing negotiations are often complicated, strategic, and multi-lateral ordeals. It is therefore unlikely that smart contracts would simplify, or even be feasible for, many forms of patent licensing.

²¹ E.g., see <https://www.uspto.gov/learning-and-resources/newsletter/inventors-eye/licensing-options> and <https://www3.wipo.int/wipogreen/en/> (both last accessed 10/12/2020).

²² Markets for patent transfers and licenses in a world with DLT-based patent records would inevitably be set up by third parties in any case.

²³ In this example, the contract would constitute something akin to a subscription.

²⁴ As licenses do not come with exclusions (and therefore enforcement) rights, there is no sensible reason for the identity of the licensee to be listed, however.

²⁵ A standard format for VPMs would be ideal for this purpose (de Rassenfosse & Higham, 2020); currently, there are no formatting requirements.

also find the associated products listed in the patent metadata. This change would result in increased levels of useful disclosure of the existence of patented information and provide a starting point for mitigating the uncertainty associated with patent thickets (Fromer, 2016; Shapiro, 2000). To incentivise patentees to use VPMS in this way, they could act as evidence of commercialisation or other usage (Trimble, 2016; Von-neuman, 1956) which may be rewarded in the form of discounted maintenance fees, a stronger case in infringement litigation, or other benefits. One particular use case is the use of DLT-based VPMS for proving version history of patent marking websites—there is currently no requirement to keep these records, but they could be crucial when attempting to claim damages on a marked item in a litigation case.

4. DLTs for structural change

We will now move on to the structural changes that DLT-based records may facilitate. Up to this point, the identified benefits of switching to a DLT-based patent records system are most useful after a patent has been granted and primarily relate to clarity in patent rights. Changes that increase clarity and ease of management may on their own provide enough justification for such a switch. This section discusses how DLTs may enable some changes that may solve more inherent problems in patent systems that have been identified in the past (see, e.g., Bessen & Meurer, 2008; Jaffe & Lerner, 2004).

4.1. Patent application

The first step in the patenting process is the filing of an application to the patent office. This step is perhaps the easiest to envision as a transaction recorded on a DLT, as described in Section 3.1. After all, the patent office already records applications and the information contained therein. Therefore, this section addresses the appropriate extent of disclosure at this early point in the process, and examines how this step would be facilitated and implemented if patent offices were to record applications on a DLT.

4.1.1. The timing and nature of invention disclosure

The status quo in many patent offices is the publication of all patent applications after 18 months (if an accelerated publication is not requested). It is not obvious, however, that there *should* be a delay in the publication of the application from a social welfare viewpoint. Disclosure is one of the most fundamental aspects of the patent system; indeed it is one half of ‘grand bargain’ that society makes with innovators (Graham & Hegde, 2015), and it is only as useful as it is timely (United States Senate, 1992). While the well-known reward theory of patents primarily focuses on the incentives that encourage inventors to invent, with the presumption that these inventions will eventually benefit everyone, there are also many benefits associated with invention disclosure. Appendix B goes into further detail about the disclosure function of patents that partially motivates our suggestion below.

We propose that there be no secrecy period automatically endowed to patent applications. If assignees would like to keep their applications secret, they can pay a fee for this privilege.²⁶ This way, secrecy benefits the public to some extent since the patent office would be better resourced as a result, while also reducing the potential impact on the prevalence of breakthrough inventions—without the secrecy option,

²⁶ The arguments against this proposition generally concern with time-to-market. For example, firms would lose some of the lead time advantage bestowed upon them by the current secrecy period, which may particularly affect small firms and technologies that take a long time to commercialise. This of particular concern if it is possible to imitate the inventions quickly in jurisdictions with weaker IP enforcement. We believe that the benefits outlined below outweigh these costs, which may, in any case, be mitigated through tools such as price discrimination dependent on firm or technology type.

high-risk R&D projects in particular may be disincentivised at the margin. In the context of DLTs, ‘dynamic secrecy’ may be facilitated by a smart contract or off-ledger mechanism that decrypts and publishes the patent application at a time of the firm’s choosing, perhaps retaining the 18-month secrecy period as a maximum. The cost of entering this arrangement would scale with the requested secrecy period.

This proposal pertains to the disclosure of the invention itself. However, there is additional information, sometimes referred to as *non-technical* disclosure (Anderson, 2016), that may be leveraged to make increased levels of disclosure more palatable to patentees at the cost of the (initial) transparency described above. To illustrate one such lever, let us assume that the utility of disclosure generally does not depend on *whose* invention is being disclosed. For this reason, it is not strictly necessary to disclose any metadata that explicitly identifies the assignee, inventor, attorney, or any other entity related to the invention until the patent is granted. This mechanism could assist in maintaining some balance between disclosure and incentives to invent, as disclosing the details of the invention would not signal the research direction of any particular applicant, merely that it is *someone’s* research direction. Further, the initial separation of invention data from inventor data is line with more general data transparency principles (Janssen, Matheus, Longo, & Weerakkody, 2017). Note that in many jurisdictions patents are reviewed for national security purposes, and may be placed under a secrecy order as a result. This process is clearly an important one, and we briefly discuss it in the current context in Appendix C.

These proposals, taken together, would have several profound effects, both on the way the patent system functions internally, and how it is utilised. First, the suggested timing of pre-grant publication will lead to earlier disclosure of new inventions, more patent office funding, or a combination of both. The earlier disclosure of inventions would not only bring with it all the benefits of disclosure, such as more timely follow-on innovation (Furman, Nagler, & Watzinger, 2018) and reduced risk of duplicative efforts (Lück, Balsmeier, Seliger, & Fleming, 2020), but may also increase the quality of patent applications.²⁷ After all, potential applicants would be less willing to disclose the results of costly R&D if they were not sure about the patentability of these results, especially if patent examination was more rigorous than at present (see Section 4.2) to combat potentially increased levels of obfuscation of invention details. A firm that is indeed unsure about the patentability of an invention in its current state (assuming a patent is necessary) would have two options: develop the invention further to ensure patentability (Cotropia, 2009 explores this option in a related context), or pay for secrecy—both of these are socially desirable outcomes. Further, recent changes to trade secret law and the emergence of third-parties specialising in DLT-based trade secret management have made early disclosure more feasible than ever—we discuss these developments in Appendix D.

Firms can currently request early pre-grant publication in some jurisdictions, and the benefits of this early disclosure are particularly pertinent for small firms who do not have the resources to commercialise a promising invention. Patent applications, and patents themselves, are vital to these firms as demonstrations of their capabilities and can be used as a signal to investors (Conti, Thursby, & Thursby, 2013; de Rassenfosse, 2012; Farre-Mensa, Hegde, & Ljungqvist, 2020; Hall, 2019; Häussler, Harhoff, & Müller, 2012; Hsu & Ziedonis, 2008; Lemley, 2000; Long, 2002; Veer & Jell, 2012). Immediate disclosure would speed up these fundraising processes, including pre-grant licensing (Hegde & Luo, 2017) and loans (Saidi & Zaldokas, 2020), and therefore the realisation of the public and private benefits stemming from these inventions.

4.1.2. Incentivising higher quality applications

Surveys consistently find that lead time and secrecy are more effective appropriability mechanisms than patents across many industries,

²⁷ Early disclosure has been studied extensively in both the theoretical (e.g., Aoki & Spiegel, 2009) and empirical literature (e.g., Graham & Hegde, 2015).

countries, and time, (Arundel, 2001; Blind, Edler, Frietsch, & Schmoch, 2006; Brouwer & Kleinknecht, 1999; Capponi, Criscuolo, Martinelli, & Nuvolari, 2019; Cohen, Goto, Nagata, Nelson, & Walsh, 2002; Cohen, Nelson, & Walsh, 2000; Harabi, 1995; Levin et al., 1987). Despite this evidence, a culture of ‘file early, file often’ is still prevalent, and indeed encouraged, as inventors can claim priority on inventions before they know much about potential utility, paths towards implementation, or even commercial viability (Cotropia, 2009). This practice has significant ramifications for the quality of disclosure—at the time of early filing, the inventor, by definition, has not refined the invention and therefore often does not have enough information to adequately delimit the scope of their patent or fulfil the enablement requirement (Cotropia, 2009). Applicants can further exacerbate this problem if they choose to merely fabricate the potential applications of the invention as part of the specification.²⁸ These applications do not need evidence to demonstrate viability (Freilich, 2019; Freilich & Ouellette, 2019). Immediate disclosure may shift firms’ incentives to ensure they have sufficient lead time to offset the risk that the patent is narrowed considerably during examination (or rejected altogether).

There exists a risk, however, that firms would make their applications extremely broad or technically obscure in response to an immediate disclosure policy, which would substitute for the desired increase in lead time before filing a patent application. Of course, the option to pay for secrecy as proposed in this paper would discourage this behaviour for valuable inventions (assuming firms are more likely to pay for secrecy in these cases); however, the patent value distribution is known to be highly skewed (Gambardella, Harhoff, & Verspagen, 2008; Moore, 2005; Pakes, 1986) and so this payment may not be worthwhile for most patents.

There is clearly a need for balance, and we do not present here any DLT-based solutions to mitigate intentional obscuring of invention specifications. As such, it is important to note that both early filing and the lack of clarity in patent applications are well-studied problems. Appendix E discusses some potential solutions to these issues that may provide the balance necessary for immediate disclosure to be a sensible modification to patenting procedure.

4.2. Examination

4.2.1. The problem of weak patents

Recent years have seen increasing concern about the quality of granted patents (de Rassenfosse et al., 2020; Jaffe & Lerner, 2004; Wagner, 2008). Much of this concern has been focused on the perceived proliferation of so-called ‘weak’ patents (Choi, 2005; Farrell & Shapiro, 2008; Lemley & Shapiro, 2005) that are likely to be found invalid if challenged in court. After all, the exclusion rights obtained by patenting an invention only exist insofar as they are enforceable. In the parlance of Lemley and Shapiro (2005), uncertain enforceability of a patent is equivalent to probabilistic patent rights—a patent gives the owner a *chance* to exclude others from this technological space. An increase in the number of invalid patents that are granted corresponds to this chance becoming smaller. In turn, the incentives to partake in risky inventive activities become weaker. These likely invalid or ‘weak’ patents impact innovation in some industries more than others (Allison & Lemley, 1998; Bessen & Meurer, 2013); and it is conceivable that more patent litigation due to the existence of these patents in a few industries could significantly increase the level of defensive patenting among large firms in these industries (Chien, 2010; Hall & Ziedonis, 2001), adding to the backlog (and grant lag) for all prospective patentees.

The reasons for the granting of invalid patents are also well-studied. In a recent work, Henkel and Zischka (2019) find that one of the primary causes of the invalidity of German patents, as determined by the German court system, is incomplete prior art search. This finding is consistent

with previous results using USPTO data (Frakes & Wasserman, 2017) which indicate that less examination time results in patents with less prior art cited, an increased grant rate, and a higher likelihood of being found invalid if challenged. It is the job of the patent office to grant *valid* patents and, in many jurisdictions, restrictions on examination times are clearly too strict for the examiners to fulfil this requirement for all patent applications. After all, if examiners are allotted less time to examine each patent application, they have less time for an extensive search of the prior art, and so are less likely to reject an application based on this prior art. At the USPTO, for example, the burden is with the examiner to show why the patent *shouldn’t* be granted. In this scenario, less examination time leads to less opportunity to find a good reason to reject an application and, ultimately, results in a higher likelihood of grant for any particular patent.

However, restrictions on patent examination time is sometimes a necessary evil. Continuing with the USPTO as an example, some seemingly innocuous changes made to the patent system in the 1980s (Jaffe & Lerner, 2004) has led to the number of patent applications to the office ‘exploding’ in the past 30 years (Hall, 2004; Strandburg, Csardi, Tobochnik, & Erdi, 2006), and the USPTO has generally not managed to keep up with this demand (Ackerman, 2011). While all patenting firms suffer the increase in grant lags associated with the resultant backlog, it particularly impacts small firms and startups. For example, there is evidence that the beneficial side-effects of patent grants for these entities are largely diminished by long grant lags—the delayed access to financing made possible by patent ownership can prove as damaging as the rejection of the application altogether (Farre-Mensa et al., 2020). While the USPTO announced “specific actions [they] have taken to help reduce the backlog” in 2011,²⁹ and the backlog has indeed decreased, the *percentage* of applications that are granted has risen at the same time (Cotropia, Quillen Jr, & Webster, 2013). Therefore, the net effect of decreasing examination time allotments may simply be a reduction in grant lags at the expense of increasing levels of litigation when these weaker, expedited patents are challenged in court—arguably an even greater source of inefficiency in the innovation system (Frakes & Wasserman, 2019).

One potential, and perhaps obvious, solution to this particular problem is an increase in patent office resourcing to allow examiners more time to consider the patentability of applications. Assuming that this is either not feasible or would not work (Lemley, 2011), we will discuss the concept of *prior art bounties* as well as the contracting-out of patent examination services for the remainder of this section, and how DLTs may enable these developments.

4.2.2. Prior art bounties

The idea behind prior-art bounties is simple: pay third parties for relevant prior art that leads to a rejection of a patent application.³⁰ This concept was first proposed and explored rigorously by Thomas (2001). The need for such a system has been echoed more recently (Ghafele, Gibert, & DiGiammarino, 2011), and even trialed in several countries, albeit in the absence of financial rewards. These trials, in the end, were a series of short-lived projects starting in the United States as *Peer-to-Patent*,³¹ and assumed that the benefits to the third-party experts would

²⁹ <https://www.uspto.gov/learning-and-resources/newsletter/inventors-eye/uspto-year-review-and-look-forward>

³⁰ Note that this proposal is distinct from outsourcing prior art searches as practiced in some jurisdictions (Yamauchi & Nagaoka, 2015) and discussed in Section 4.2.3. However, both outsourcing and the current proposal serve a common purpose: to widen the scope of the prior art search and attempt to optimise patent quality with limited examiner resources.

³¹ The results of these programs likely inspired new or updated implementations of third-party submission processes around the world, including those implemented in the United States (112th Congress of the United States of America, 2011), Australia (IP Australia, 2018), and at the World Intellectual Property Office (WIPO, 2012).

²⁸ This practice is legal and done in part to illustrate invention utility and preempt potential infringements

outweigh the time cost of the prior art search. These ‘strong’ motives (Noveck, 2006) included the “opportunity to produce better patent quality”, as well professional motivations such as the will to keep their domain knowledge up to date and gain recognition among their peers and future employers. We argue that these motivations may not be sufficient for sustained participation of experts in such a program (Wijnhoven, Ehrenhard, & Kuhn, 2015). Prior-art bounties take a more pragmatic approach by offering a financial reward alongside these other incentives. The aim of this approach is not to encourage third-parties to conduct in-depth searches in the hope of finding relevant prior art. The motivation is to give examiners access to large, preexisting reservoirs of expert knowledge that can quickly identify relevant prior art—thus making the job of examiners easier, improving patent quality, and financially rewarding experts for their vigilance.³² It is important to note that the following suggestions do not aim to replace the entire prior art search process, but merely assist examiners in making patentability decisions that would, hopefully, result in fewer validity disputes.³³

A prior-art bounty system would consist of two stages: a submission stage and a review stage. Submissions begin as soon as the patent application is made public and close while the application is under examination. The submission may cost the ‘bounty hunters’ some nominal fee per submission to disincentivise bad actors and minimise the number of low-quality submissions. All submissions would be published on the DLT and, therefore, be publicly viewable. The public and the applicant are thus able to keep track of submissions.

The examiner, during their prior art search, first considers the submitted art for each patent application in order of submission.³⁴ If they deem a particular piece of submitted prior art relevant and use it as justification of a rejection notice of any kind, the bounty hunter receives a reward or *bounty*. At this point, fees from all other submissions would not be refunded, even if the examiner did not see the submission before the rejection—this mechanism incentivises the submission of prior art to applications where the hunters do not believe any relevant prior art has yet been submitted. The benefits of such a bounty system are clear: third parties are incentivised, professionally and financially, to provide relevant prior art to the patent office.³⁵ These incentives may be further enhanced when a patent is of particularly high value, as competitors would like to see the scope of these patents narrowed as much as possible. Competitors may know more about the technology and state of the art than do the examiners. Therefore, in a system that restricts examiners in their ability to conduct an effective prior art search, prior-art bounties appear to us as a way to reduce this load, reduce the incidence of weak patent grants, and ultimately work to ensure the integrity of granted patents.

The administration of a prior-art bounty system need not be resource-intensive—many of the processes involved can be automated, on or off the DLT. As an on-DLT example, the submission of prior art and associated submission fee could append the submission to the patent application metadata and initiate a smart contract between the hunter and the patent office. This code would monitor events on the DLT and execute appropriately, dependent on the fate of the application. If the

patent is rejected, the contract would check the prior art used in the rejection justification (which is also recorded on the DLT). If it matches the prior art submitted by the hunter, then the hunter is paid their dues. Else, if the application is withdrawn or abandoned before examination, the fee is refunded. Else, the patent is examined and either rejected or allowed with no reference to the submitted art, and the fee is forfeit. Every transaction from the time the prior art is submitted is automated and requires no human input beyond the usual actions of the applicant and examiner. Acceptance conditions of this smart contract could also be set automatically, taking all public information about the application into account and making an offer to potential bounty hunters—applications for notoriously marginal or technical types of inventions may have their fee-to-bounty ratios and the absolute levels of these quantities adjusted accordingly.³⁶ The precise determination of the financial reward is out of the scope of this work, but we will note that it should consider total savings made by both patent offices (in terms of examiner time) and industry (in terms of invalid patent grants). Further, while it is difficult to judge the *ex-ante* value of a patent, it will still take time to examine. As such, even if a patent is ‘worthless’ (Moore, 2005), it is still worth financially incentivising prior art submissions, and particularly so if low-value patents are less likely to be valid. For the same reason, rewards should not depend on applicant characteristics, such as firm size. In any case, the optimal reward structure of a bounty system is an interesting theoretical question that could be addressed in future work, perhaps from a game-theoretic perspective.

Prior art bounties come with some issues that need to be addressed in detail. Some of these are very jurisdiction-specific; we discuss some examples of these in Appendix F. Others are more general, the most pressing of which concern the potential for collusion between the different parties involved. Of course, there is a need to ensure that the incentives structures of bounty hunters, applicants, and examiners are appropriately designed. Safeguards in the form of patent office oversight, legislation, or pricing need to be put in place to ensure that there is no collusion between these parties or anonymised foul-play that can undermine the value of a bounty system. As a straightforward example, there is nothing in this system preventing the *examiner* from submitting prior art as an anonymous bounty hunter and then using their submission as justification for rejection. Another problematic scenario is one in which the people or firm doing the prior art search before filing the application may leave an important piece of art out of the application, knowing that they would be able to submit this as a bounty hunter and claim the bounty, while potentially being hired to amend the application in light of this prior art. This concern is more acute if the submitted documents have not been timestamped by an authority or DLT-record (see Appendix G).

One way to mitigate this kind of behaviour, in addition to outright punishment, is to limit submissions to the non-patent literature only. After all, these references are generally the most challenging and time-consuming for examiners to find (Sampat, 2005; Thomas, 2001). This kind of prior art has the additional benefit of being less intimidating to those who are much more familiar with a field-specific set of documents (e.g., academic literature) than the examiners or even the applicants themselves. This limitation divides the search effort more efficiently, as patent examiners are likely more proficient at searching for relevant patents than they are at searching for non-patent literature.³⁷ Another

³² A recent review by the US Government Accountability Office found that only 42% of USPTO examiners ‘always’ or ‘often’ examine applications concerning subject matter of which they have some knowledge of the prior art (GAO (2016)).

³³ GAO (2016) reports that while only 17% of USPTO examiners had seen third-party submissions, in cases where submissions were made, they were used to reject a patent application 20–25% of the time.

³⁴ This queue-based prior art check is already standard practice, with applicant-supplied prior art reviewed first before an independent search by the examiner.

³⁵ One can imagine this process being streamlined for large organisations employing many experts, whereby the experts submit relevant prior art to their legal team or technology transfer office for formatting and submission, with the organisation pooling and sharing expenses and bounties.

³⁶ This programme could, initially, take a form similar to that of Peer-to-Patent, where applicants could opt-in to such a third-party submission system for particularly technical subject matter.

³⁷ A recent work, however, has examined the types of prior art used by the United States’ Patent Trial and Appeal Board (PTAB) to justify the invalidation of a granted patent after an Inter-Partes Review (Yelderman, 2019). The authors found that while most obviousness invalidations cited a U.S. patent (76%), almost half (47%) cited a printed publication such as a book or journal article (note that multiple citations are allowed).

foul-play mitigation strategy is to ensure that bounty hunters are subject to a robust identity verification procedure during the bounty hunter registration process. These potential abuses of a bounty system suggest a new administrative role for patent offices to weed out bad actors—an administrative cost that must be weighed against any change in granted patent quality expected as a result of prior art submitted through this system.

4.2.3. Contracting out examination

Examination of patent applications may be considered to be one of the main privileges of patent offices. However, in the face of concerns about the quality of granted patents, restrictions on examination time, and skewed incentive structures in favour of issuance, we may be forgiven for thinking that there may be alternative procedures that can do better. In this section, we propose the contracting out of examination services,³⁸ and how DLTs facilitate the implementation of such a system that may otherwise be considered too complicated or trust-dependent to administer.

Many patent offices are self-funded by applicant fees, with the burden on the examiners to prove invalidity. Additionally, patent offices have a monopoly on patent examination services (Abramowicz & Duffy, 2008) and can thus perform poorly with little risk (Domberger & Jensen, 1997). In combination, these facts encourage the grant of as many patents as possible while skewing incentives away from ensuring high-quality examination. Third-party contractors would be incentivised to conduct a high-quality examination—they would not only be subject to patent-office oversight but would also have to compete for examination contracts. This competition would also incentivise contractors to find ways to examine patents as efficiently as possible, and spur on technological innovations that the patent office lacks the resources and incentives to develop.³⁹

The contracting out of government services is not a new idea, and the piece-wise nature of patent examination makes this a particularly suitable application of this process—the prior art search or substantive examination of a single patent application may constitute a discrete contract. Blueprints for contracting out prior-art search to qualified third parties already exist in multiple jurisdictions (Kazenske, 2003; United States Patent and Trademark Office, 2003; Yamauchi & Nagaoka, 2015). In some jurisdictions, patent examiners often leave to become patent attorneys (Drahos, 2010; Tabakovic & Wollmann, 2018) and are, therefore, very familiar with the patent examination process, perhaps more experienced than many of the examiners at the patent office. This attrition suggests that there exists a pool of examiners, some of which may be willing to come back to this role as a private practitioner, with more flexible hours and higher pay. On the other hand, examiners that remain would be able to spend much more time on marginal or challenging cases, which may be appealing to many technically-minded people in want of stable employment.

Prior art search may justifiably be considered an integral (and not separable) part of the examination process (Jeffery, 2002). That is, integral to a high-quality examination is a high-quality prior art search. If competition for examination contracts is fierce, therefore, third-party examination would inevitably lead to innovations that change how prior art is searched for, at least compared to the way it is done in the context of examination within the patent office. This effect may be further amplified when all transaction data, including all third-party prior art submissions and their success, are publicly available and stored on a DLT (see Section 5.2). The argument for such a system is strengthened further when considering a DLT shared between multiple offices (see Section 5.3), who can access the same prior art information

for whole families of patents and applications, reducing duplicative search efforts.

In Japan, outsourcing the prior art search component of patent examination is already practiced. Yamauchi and Nagaoka (2015) observe that, when given a choice, Japanese patent examiners were more likely to choose to outsource prior art search for less complex technologies. The examined applications that had the search outsourced had significantly fewer appeals against both rejection and grant decisions, and final decisions were reached more quickly when compared to similarly complex applications for which search was conducted internally.⁴⁰ This evidence suggests that, at least in the case of less complex technologies, outsourcing prior art search is effective at both increasing the quality of examination and reducing pendency times; the potential to extend these effects to more complex technologies could be realised by a system where third-party examiners have the time and technology to integrate these effective searches into the examination process. Such a system would require significant oversight by patent offices to ensure the integrity and robustness of the examination process, as discussed further in Sections 4.2.3 and 5.1.

4.2.3.1. Small-scale experimentation. We suggest that the technical components of the current proposal could be trialed *within current office procedures*, complementary to current systems. This experimentation may be conducted by simply recording all examiner activities and correspondence on the DLT. After this technical foundation is put in place, an initial opt-in system could be designed to take up patent office slack and could be restricted to particularly non-controversial or non-complex technology fields. Using the outsourcing of prior art search by the Japanese Patent Office as a model (Yamauchi & Nagaoka, 2015), examiners themselves could have discretion as to which patents can easily be examined by third parties. Further testing of third-party examination may even involve other national offices that may be better resourced or have particular areas of technological specialty. Such contracting out is already being done by some small patent offices or collections of such offices (albeit in the absence of a DLT). One example is the Support System for Patent Applications Management for the Central American countries and the Dominican Republic (CADOPAT), which also extends to the African Regional Intellectual Property Organisation.⁴¹

4.2.3.2. Incorporating review processes. DLT-based examination would also facilitate other permutations of public-private cooperation in patent examination. For instance, it is possible to set up a system to help ensure the incentives of private examiners are aligned with the patent office's, such as the second-pair-of-eyes review (SPER) program (Allison & Hunter, 2006). In the current context, this process would be a request for the patent office to re-examine a patent application after the private examination is complete. Mark Lemley and others have also suggested that a patent surviving reexamination should result in an earned presumption of validity, should its validity be challenged (Lemley, 2011; Lemley, Lichtman, & Sampat, 2005; Lichtman & Lemley, 2007). These ideas are discussed in more detail in relation to private examination in Appendix H.

Contracting out examination services not only provides a natural application of SPER as a way to 'gold-plate' patents—first by going through a third-party examiner, then through a patent office examiner—but also provides checks on the performance and reliability of third party examiners, perhaps in addition to random checks (see also Abramowicz & Duffy, 2008). This system also solves, in part, an important problem with private examination: the potential for collusion between the applicant and the private examiner. At present, while the

³⁸ Note that this concept has been discussed previously in the legal literature (Abramowicz & Duffy, 2008).

³⁹ This innovation is also likely to occur outside of the contracting firms by specialised service providers.

⁴⁰ There are indications that these effects are driven in part by a broader scope in the prior art search of the contractors.

⁴¹ https://www.wipo.int/meetings/en/doc_details.jsp?doc_id=345737 (last accessed 10/12/2020)

incentives, resource constraints, and protocols at some patent offices are undoubtedly skewed in favour of granting patents (Frakes & Wasserman, 2015; Schuett, 2013), examiners are not incentivised to favour some ‘customers’ over others. When examination is private, there is much more room for these activities to occur, and less direct oversight to detect such activities. Granted, all official communications, including all justifications for changes to the applications, would be recorded on the DLT for all to see, but this recording does nothing to prevent, for example, leniency towards marginal patents in exchange for return business. Of course, one way to mitigate this effect is to impose penalties on contractors where evidence of misbehavior is uncovered. As the entire examination process is public due to being recorded on the DLT, firms would also be able to act as watchdogs during the examination of their competitors’ patent applications, as they too have a stake in the outcome.

However, SPER may also provide a way to mitigate this kind of collusion. We foresee SPER being used in three circumstances:

- When an applicant wants added certainty in their patent rights via a presumption of validity should the patent be litigated;
- When a third party believes a patent has been wrongly allowed; or
- When an applicant believes their application has been wrongly denied.

In each of these cases, the party initiating the SPER pays for the examination to start the process. This deposit guarantees that this party is financially committed to the outcome. However, if the patent office does not agree with the result of the private examination, then the *contractor* must pay the examination fee and forfeit their initial examination fee (which was recorded in the DLT), while the initiating party is refunded. Assuming the SPER process is only utilised for valuable patents (from the point of view of either the applicant or a competitor), this mechanism would assist in incentivising private examiners to try to be at least as strict and thorough as a patent office examiner during their examination process, and particularly so when the patent protection sought is valuable (to anyone). The above integrates the incentive schemes for private examination explored in Abramowicz and Duffy (2008) with the SPER system (Allison & Hunter, 2006; Lemley, 2011), while DLTs record and make transparent, in real-time, the whole process in a way that relies on no single entity to do so.

4.3. Keeping patents in force

After grant, patentees must pay maintenance fees to the patent office to keep a patent in force. From the perspective of the patent office, the monitoring of maintenance fees via a smart contract has the potential to reduce administrative costs associated with keeping patents in force. DLTs may also enable a schedule of fees that is more optimal from a social welfare viewpoint.

A simple implementation could manifest itself as follows: upon patent grant, a smart contract would be initiated. It would track the receipt of associated maintenance fees and the time at which this transaction is added to the DLT. Another smart contract is initiated regarding the next maintenance fee event. If the patentee does not pay the fee, then the patent is considered expired, and the invention enters the public domain.

Such a pay-as-you-go maintenance fee system may function as follows: inventions would automatically obtain some minimum period of protection on receipt of the patent issue fee (on the order of a few years), during which time the patent owners may pay their maintenance fee at any point. Three factors determine the amount paid: time since grant; additional time requested; and time until expiration.

Intuitively, the fee should increase with time since grant and the additional time requested, and decrease with the length of time until expiration to offset the risk that firms take on by paying maintenance

fees early.⁴² Of course, the maximum patent term should still be limited in time, and the invention would fall into the public domain at the latest at this point as usual. This kind of scheme allows patentees to pay maintenance fees commensurate with the present value of the invention to them and rewards them for risk-taking when this associated future income stream is uncertain. A DLT-based maintenance fee system would effectively remove administrative burden from this process altogether; a smart contract could, knowing the time of grant and the expiration date immediately prior to payment, be able to calculate the additional time requested directly from the amount paid and apply this extension to the patent metadata after confirmation from the patentee.

As a first step, this pay-as-you-go maintenance fee system can easily be matched to the current pricing model. For example, the parameters that determine the relationship between the three factors listed above and the resultant fee can be calibrated such that the price of a set expiration time extension (e.g., one year at the EPO and four years at the USPTO) at a particular time match the current prices for the same extension. The flexible pricing described can then be entirely optional for those firms who would like maximum control—after all, this flexibility would integrate nicely with automated payment level optimisation based on, for example, the present value of the patent rights, the uncertainty of this future income, or other technology performance metrics (Jin et al., 2011). Other firms may prefer some predefined payment schedule, such as those currently in place. For these firms, the maintenance fee system would not change unless they would like to deviate from it.

5. New roles for patent offices

Patents represent a government-granted right. As such, patent offices, as representatives of their respective governments, have a vital role to play in every proposal contained here. This fact remains true if the decentralisation of various patent-related processes is taken to its most extreme. Here we will go into detail about the new roles that the patent offices would need to assume to achieve the kind of flexibility and efficiency gains that are intended by the proposals above.

5.1. Oversight and regulation of patent examination

Supposing patent examination may be contracted out in a piecemeal manner, contracts need to be distributed and monitored in a way that minimises the risk of collusion between third-party examiners and applicants. Here, we will remain agnostic to the choice of contract allocation mechanism.⁴³ We will, therefore, focus on patent office oversight of third-party examination services and the challenges that may arise independent of a particular contract allocation system.

It is vital to ensure complete transparency of the examination process if a third party carries it out. To this end, all communication between the applicant, the patent office, and the contractor should be recorded on the DLT in a standard format. Patent offices would therefore be able to conduct significant oversight, some of which may itself be outsourced to third parties specialising in, for example, anomaly detection or

⁴² There exists a large body of literature, both theoretical and empirical, that examines patent fee structures (both before and after grant) and the incentives that different structures produce (Baudry & Dumont, 2006; Cornelli & Schankerman, 1999; de Rassenfosse & van Pottelsberghe de la Potterie, 2013; Gans et al., 2004; Scotchmer, 1999).

⁴³ It may be useful to keep one’s favoured mechanism in mind, however. Classes of allocation mechanism include: applicant’s choice of private examiner, patent-office-fixed contracts for individual patents or bundles offered to trusted firms, and patent-office-moderated auctions of the same. Catalini and Gans (2016) gives an example of how the latter may function via a DLT in a generic auction context.

examination quality monitoring.⁴⁴ Furthermore, contractors would need to declare any conflicts of interest they may have before accepting the contract for any given examination. It is crucial to ensure not only that the grant incentives of the examiners are orthogonal to the incentives of the applicant to mitigate collusion, but that both parties benefit in the case of the grant of a high-quality patent. That is, examiners must benefit when they grant a high-quality patent, and for reasons unrelated to those benefiting the applicant. At the same time, examiners must not be disadvantaged due to a well-reasoned rejection.

The DLT-based record described above could then be aggregated by contractor and would contain all information of interest to potential applicants, including the cost of past examinations, technological specialties, and examination times. This aggregate report could also contain post-grant information, such as any invalidity rulings or other litigation outcomes related to these patents. The patent office may also use this information to optimise their monitoring processes; for example, if patent office examiners conduct random checks on examined applications (Abramowicz & Duffy, 2008), then the probability of selecting a patent examined by a particular contractor could be determined by their record through a publicly viewable reputation metric.⁴⁵

To become a registered examiner in the first place, a training program similar to that undertaken by patent office examiners should be required, and these contractors should have attained (at minimum) the same qualifications as the equivalent patent office examiners. To ensure that their examination skills remain sharp, contractors should have to examine some minimum number of applications every year. Furthermore, to ensure that their knowledge is kept up-to-date, they should enroll in continuing education classes provided by the patent office. Additionally, technology-specific examination qualifications can be provided by the patent office in order to train both in-house examiners and contractors to assist in the examination of applications in legally or technically challenging fields. The registration of third-party examiners may be suspended or terminated by the patent office in the case of particularly poor performance, or malfeasance.

The primary function of patent offices is to examine patents and, therefore, it is during this pre-grant period that most patent office expense is incurred. However, to lower the upfront cost to applicants, these expenses are subsidised by maintenance fees paid by patentees. Additionally, current patent fee schedules at some offices differentiate between firms of different sizes: smaller applicants pay less at each step of the process. These details are important in the context of third-party examination, as contractors would need to be paid in full for their work, thereby potentially increasing upfront costs for the applicant. Therefore, regardless of the contract allocation mechanism, there exists a question about how (and how much) the patent office would cross-subsidise these examination services should such a system materialize. The answer is sensitive to the particular implementation of this system, so we leave this question open and amenable to theoretical scrutiny (See, e.g., Gans, King, & Lampe, 2004).⁴⁶

5.2. Patent data organisation

One consequence of recording patent-related transactions to a DLT is that this information is automatically preserved, immutable, and all in

one place. Exactly how particular types of information is to be stored will not be discussed here⁴⁷; however, note that it is currently standard practice to store files in a dedicated database with any metadata, hashes, and links to the file location stored on the DLT itself (Xu et al., 2017).

Disclosure of inventions is not only a function of the clarity of the invention description but also of accessibility. With all patent-related data in one place (or at the very least all metadata with permalinks to full files), third parties, or even patent offices themselves could start building flexible tools for the public to access this information easily. These data should include all files relevant to each granted patent in a standardised format,⁴⁸ including datasets, CAD files for 3D models, and more specialised formats for, e.g., molecular structures. Access to these data for large-scale analyses by commercial ventures could be offered for a fee.

This single source of information could also facilitate the examination of new patent applications, through the development of advanced search techniques exploiting it. If examination is done by the patent office, then these tools can be developed in-house or in partnership with a third party. However, if examination can be contracted out, then the competitive environment created would likely drive innovation in this direction at a much faster pace than what can be achieved by the patent office alone. Software designed to examine novelty and non-obviousness would be able to use the data contained in all of the data types listed above (alongside the patent description itself) to facilitate these judgments and increase both the scope and nature of prior art search of patent documents and applications.

Changes to this database would be, by definition, events published on the DLT. These events, being publicly accessible in real-time, could trigger smart contracts running on different ledgers (of any kind) or local programs on a centralised system. This event-by-event information pipeline is much more conducive to these kinds of activities than the APIs currently available, which require requests for specific information—the programs would detect relevant new events almost immediately, and be able to react automatically. This change may open up new possibilities for technology management and monitoring for anyone interested and at the very least would significantly increase the timeliness and efficiency of current efforts to do so.

5.3. International coordination

One obvious question that arises when discussing distributed ledgers is: among whom is the ledger distributed? This question is important to address, as DLTs are often seen as a way to decentralise authority and control of a system. However, there are other justifications for the use of DLTs. Some firms use private blockchains, for example, to ensure all transactions of a particular type are recorded in one place, whether it be supply chain management (Kshetri, 2018; Saberi, Kouhizadeh, Sarkis, & Shen, 2019), or work-flow management (Fridgen, Radszuwill, Urbach, & Utz, 2018). These applications of DLTs are centralised by design—the DLT is used as a tool to assist in the recording and management of many actions by many different actors within (or contracting for) an organisation. The incentive to ensure an accurate timeline of transactions is simply the utility gained by an immutable and centralised ledger, which can facilitate oversight and information aggregation for functions such as performance monitoring and third-party audits (Dai & Vasarhelyi, 2017).

However, many patent offices around the world may find utility in a DLT for their own applications. In the case where many organisations would appreciate access to the same DLT, a consortium (see Appendix A)

⁴⁴ Outsourcing in this manner is not unusual. For example, financial auditing is outsourced to large accounting firms in many countries.

⁴⁵ These kinds of checks for examined applications are currently not possible at patent offices due to resource constraints (Allison & Hunter, 2006).

⁴⁶ There are good arguments for raising the necessary funds through sensible fee structures, both pre- and post-examination (the latter is examined in more detail above in Section 4.3). Application fee increases at the USPTO in the 1980s acted as a filter for low-quality patents, particularly for firms with already large patent portfolios, with little effect on the quality or number of patents from small firms (de Rassenfosse & Jaffe, 2018).

⁴⁷ See, e.g., Benet (2014) for an example of a popular DLT-based file storage system.

⁴⁸ There is no discernible benefit to making *all* files public on application that would outweigh the disincentive for disclosure (or incentive for obfuscation) such a policy would create.

could provide the basis for a consensus mechanism for the DLT and provide some assurance that the ledger is indeed immutable (or very close to it).⁴⁹

DLT consortia, in particular, have been experimented with for a number of years within certain industries. These sectors include, but are not limited to, banking, finance, insurance, healthcare, and transport.⁵⁰ These applications could serve as models for a DLT-based patent records system, as many serve the same underlying purposes: a reduction of administration costs, increases in system efficiency, and easier data sharing and transactions between participants, all of which are also desirable for a global system. The use of a consortium of patent offices, however, has an additional benefit that is generally not desirable for these other applications: transparency. Indeed, this may prove to simplify implementation significantly, as privacy is generally not a concern where patents and published applications are concerned.

Most jurisdictions have their own assignee or applicant identification systems. However, there are significant benefits to harmonising these systems across countries, not least having clear and transparent ownership and relevant corporate structure information across patent families. A DLT-based global patent information database would be a good starting point for such harmonisation. Firms must elucidate ownership structures (which often stretch across jurisdictional boundaries) for other purposes such as taxation, and as such, there are few socially beneficial reasons for obscuring the identity of potential enforcers of IP with competitors and governments.

6. Barriers to adoption

It is clear that there remain many important questions that require answering before a DLT-based patent records system of any kind can be implemented on anywhere near the scale proposed in the present paper. Here, we briefly discuss some of these issues.

6.1. Legal concerns

Some of the more obvious unknowns concern the law. First, it is essential to uncover any potential conflicts between the kinds of DLT-based records described here and current laws (e.g., the Leahy-Smith America Invents Act in the United States), international treaties or agreements (e.g., the Agreement on Trade-Related Aspects of Intellectual Property Rights), or national constitutions. Further, considering a new system *without* any of the structural changes outlined in Section 4 or any conflicts with the official documents listed above, is a shift to an immutable and public record of patent information going to take a substantial amount of legislative work?

Secondly, interactions between DLT-based records and patent litigation must also be addressed. Expensive, lengthy, and excessive litigation proceedings are a well-known inefficiency within current patent systems. How might the current proposals be augmented to alleviate some of these systemic problems to make patent cases less numerous and less costly for the litigants and, directly and indirectly, society-at-large? This question naturally overlaps with other lines of enquiry concerning the future of court systems and legal processes more generally (Cabral, Chavan, Clarke, & Greacen, 2012; Remus & Levy, 2017); however, in the case of patent systems, any solutions may be able to be trialed at a smaller scale through current administrative review procedures (e.g., the Patent Trial and Appeal Board in the United States).

Lastly, patents are legal documents by nature, and patent systems

themselves are bound by laws and regulations that define the boundaries within which they operate. In many jurisdictions, some of our proposals require changes to the law that patent offices have no power to make. However, by having a DLT already in place to solve issues that do *not* require these legislative changes, implementing more ambitious changes will become more obvious and demand less political capital. Furthermore, international coordination that takes place within the framework of the European Patent Convention, the Patent Cooperation Treaty or the IP5 may provide an impetus for many of these changes. While it is important to keep in mind that patent offices are public institutions that are often beholden to the actions of legislators, they do also have some autonomy to fulfil their mandates.

6.2. Fair competition and regulatory concerns

Governments currently conduct patent examination with transparent fee schedules. Additionally, current patent office funding models use maintenance fees from granted patents to subsidise the examination costs of new patent applications, particularly for small firms. These two aspects of the current systems force two open questions regarding fee setting. First, how may we use fees paid directly to the patent office (e.g., application, issue, and maintenance) to subsidise private examination costs so that potential patentees (of all sizes) are not dissuaded from filing in the first place? Second, how can we ensure that private patent examination is priced fairly for all applicants?

Competition in examination services may force costs down, especially if there exists a database of certified examiners and their pricing. Firms active in R&D are likely to file for many patents over an extended period, so return business is also important—this may act as a quality enforcement mechanism in addition to those mentioned in Section 4.2 and help firms to ‘get what they pay for.’ However, it is clear that the private examination industry would require clear regulations to make sure that small firms are not priced out of high-quality examination services. It is also clear that these services would cost much more than current patent examination in the absence of cross-subsidisation with renewal fee income.

Another regulatory concern is to mitigate the risk of patent attorneys (representing the firm’s interest) and private examiners (representing the public’s interest) colluding to prolong examination proceedings, as they both have an incentive to extract as much money in fees as possible from applicants, who are paying both parties. Applicants are, of course, paying both parties currently; however, as the examiner receives a flat ‘reward’ per patent the incentives of these parties are actually opposed: the attorney extracts more fees via a prolonged examination, while the patent office extracts more fees by examining as many applications as possible (that is, as short an examination as possible). Additional regulation or incentive-shifting mechanisms must be put in place to make sure that examination does not take longer than necessary to assess the validity of the applications—for example by restricting the number of times an applicant can amend their application, which may have the additional effect of narrowing the scope of the average patent to avoid the additional costs of an appeal and potential final rejection.⁵¹

Implementation of a decentralised examination system also raises questions about who exactly is going to do the examining. What kind of background should an examiner have, and what kind of additional training must they go through before they become private examiners? The default answer to these questions is clearly ‘whatever we currently ask of patent examiners,’ but current examiners are quite specialised concerning the subject matter that they examine (Righi & Simcoe, 2019). This specialisation prompts further questions regarding field-specific training courses and certifications, and maintaining these

⁴⁹ A number of highly flexible infrastructures and tools aimed at building and running DLT consortia have been gaining traction recently, the most prominent being Hyperledger (Androulaki et al., 2018).

⁵⁰ See, e.g., we-trade.com, marcopolo.finance, b3i.tech, hashedhealth.com and dlt.mobi, respectively (all last accessed 10/12/2020).

⁵¹ In the literal sense of the word, rather than its somewhat misleading use by patent offices—indeed, a ‘final’ rejection can be a misnomer because an invention can never be definitively rejected in some jurisdictions.

certifications at regular intervals. One can also think of optional advanced courses that could allow examiners to command higher fees by providing higher levels of subsidisation to these more highly trained individuals, with the hope that their training results in higher-quality examination.⁵² These details are fairly minor in the grand scheme; however, they highlight the significant upfront cost of implementing such a significant change to a government-run system.

6.3. Security concerns

With many new technologies come new security concerns, particularly in the digital realm. DLTs are no exception to this; however, the particular security concerns depend almost entirely on the type of DLT that is in use, and the sensitivity of the information stored on it. In the case of patent information, transparency is generally the goal, and when secrecy is required, usual forms of encryption are available. Many personal details of parties involved in the patenting process, such as detailed contact information for inventors or attorneys, are not currently public information and there is no reason for these details to be placed on the DLT—the patent office must still retain this information centrally. Therefore, privacy is no more a concern than it is under current systems if usual precautions are taken.

The security concerns within a DLT as proposed here, therefore, are limited to two aspects: transaction validation and identity. The former depends on the method of consensus and the security of the nodes taking part in this consensus. Many existing consensus mechanisms are very robust to attack (Baliga, 2017; Cachin & Vukolić, 2017; Mingxiao, Xiaofeng, Zhe, Xiangwei, & Qijun, 2017), and the consortium approach suggested above would only allow approved nodes to take part in transaction validation in any case—an attack on this system would require actual infiltration of many of these nodes. Identity, on the other hand, is a much greater concern. To fraudulently add a transaction to a DLT (at its simplest), such as a transfer of IP, a ill-intentioned actor must acquire the victim's credentials. This act may not be very different from uncovering a password, except that this kind of attack could result in the loss of control of extremely valuable intangible assets. For this reason, any such DLT-based system would require multiple levels of authentication for certain types of transaction. The patent office has a significant role to play in the regulation of patent-related transactions, and that role may also extend to the policing of suspicious activities on the DLT and putting in place systems that ensure the validity of transactions. The use of DLTs for identity management is an active area of research (Kuperberg, 2020; Lesavre, Varin, Mell, Davidson, & Shook, 2019; Shrier, Wu, & Pentland, 2016); these may present much more secure modes of transaction than existing DLTs currently offer.

6.4. Theoretical concerns

Even small modifications to the structure and function of complex social systems, particularly those that are permeated by substantial financial incentives, may have far-reaching effects that are difficult to predict in advance. Further, some of the modifications we suggest have only become feasible to implement in recent years. For these reasons, there is both a great need and a dearth of theoretical research that attempts to model the effects of significant structural changes to patent systems.

Take private examination services, as described in Section 4.2, as an example. This set-up involves up to four parties (applicant, examiner, patent office, and decision-challenger), each with their own set of priorities and incentives. A game-theoretical analysis of this scenario may uncover strategies that entirely undermine the desired effect of this structural change. It is therefore essential that this proposal is 'incentive-

compatible' (Hurwicz, 1973).

As a second (perhaps simpler) example, pay-as-you-go maintenance fees may significantly change the way firms signal their perspective on the present value of their patents. At present and in most jurisdictions, firms pay period-by-period to keep their patents in force, which limits the inference outsiders can make about firms' own assessments of this value. When fees are flexible, a whole new set of strategies opens up that would allow firms to choose how much information about these assessments they would like to disclose through their payments. Also, on the patent office side, decisions would have to be made about the optimum parameters that govern the fee payment calculations in such a system. These calculations will depend on, for example, the expected value distribution of granted patents. Considerations would need to include the desired levels of revenue balanced against outcomes for social welfare, in more detail than required at present. On any scale, shedding light on the potential ramifications of policy changes is vital.

7. Conclusions

Taken individually, most, if not all, of the possibilities for improvement of patent systems offered in this article do not *require* that patent offices transition to a DLT-based records system. However, the flexibility unlocked by such a system could be invaluable to the trial and implementation of future modifications to innovation policy, from minor tweaks to wholesale reforms. There exist exciting proposals in the legal and economic literatures that suggest seemingly straightforward solutions to many of the issues that plague current patent systems. However, most solutions would constitute major administrative disruptions and place significant and continuous financial burdens on patent offices or their users. A DLT-based system not only makes many of these ideas administratively feasible, but can also be tested in a step-wise, scalable, and very public manner. Further, a DLT may facilitate reliable information sharing among offices and patentees around the world, reducing the burden on examiners, and perhaps even accelerating harmonisation efforts. DLTs also have additional transparency and archival attributes baked-in. A patent *should* be a privilege bestowed on those who take resource-intensive risks to explore the frontier of our technological capabilities, as a reward for their achievements. Full transparency of these rewards and the ways they are decided is very much in the public interest, as it is society-at-large that pays for inefficiencies, both administrative and economic, that exist in today's systems. DLTs can enhance this transparency.

From an organisational perspective, a DLT-based patent system has the potential to remove current bottlenecks in patent processes by making these processes more efficient, rapid, and convenient for applicants without compromising on the quality of granted patents. However, in order for such a system to meet this potential, patent offices must play important new administrative and oversight roles including: mitigating the impact of bad actors; arbitrating in cases of conflict; organising patent data and producing tools for access to this data; and leading international coordination efforts.

There exist important questions about the future of IP policy more broadly that will inform the design of a DLT-based patent system. In particular, any discussion concerning reforms that may fix inefficiencies in patent systems should start with the following question: are current patent systems, *if they were working exactly as intended*, the best way to encourage technological progress and optimise societal welfare? The answer to this question is impossible to know. After all, technological progression, and the way technology benefits society at large, is highly heterogeneous and so requires a heterogeneous set of approaches to innovation policy (Hemel & Ouellette, 2019). However, an obvious first step is allowing our innovation-inducing policies of choice to dovetail effectively by making our current systems much more amenable to modifications that may address current inefficiencies and any new ones that arise. For example, how might we build-in interoperability between patent systems and other forms of innovation-incentivising policies such

⁵² In a recent work, Righi and Simcoe (2019) observe lower grant rates for more specialised patent examiners, indicating more stringent examination.

as patent box regimes, prizes or R&D tax credits? Further, could the use of DLTs be extended to other types of formal IP such as trademarks or copyrights, or are these somehow so fundamentally different so as to warrant an entirely different framework? Answering these questions will be an endeavour stretching across disciplinary boundaries, and we hope that future research will address these as the modernisation of IP systems and policies continues in all its forms.

Author details

Gaétan de Rassenfosse is Assistant Professor at École polytechnique fédérale de Lausanne (EPFL) in Switzerland. Kyle Higham is Post-doctoral Researcher at the Institute for Innovation Research at Hitotsubashi University in Tokyo, Japan.

Funding statement

This work was supported by a grant from the Swiss National Science Foundation (grant number 10DL12_183088).

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors are grateful to Marco Barulli, Emma Francis, Alan Marco, Geoff Sadlier, Kees Schuller, Claudia Tapia, Adam Jaffe, Joachim Henkel, Dominique Guellec and Beth Webster for useful discussions, and to three anonymous reviewers for their insightful critiques and suggestions.

Appendix A. Overview of distributed ledger technologies

A distributed ledger is a list of transactions or instructions that are distributed across a network of nodes. Every entity with access to the network (that is, able to add a transaction to the ledger) is assigned a wallet, with which a private key and a public key are associated. The former is used to generate a digital signature to authorise transactions while the latter is a public identifier that can be used by others (that is, anyone who is allowed to read the ledger) to verify the authenticity of the signature and, therefore, the transaction.⁵³ Every transaction that an entity (which is not necessarily a node) wishes to have recorded on the ledger is broadcast through the network and added to each node's copy of the ledger. Through this mechanism, the network acts as a witness to a transaction (Holden & Malani, 2019). If the transaction is found to be valid, the network can come to a consensus that the transaction did indeed occur. Technologies that distribute a ledger across a network in this fashion, with an in-built consensus mechanism that ensures that all participants agree on the contents of the ledger, are known collectively as distributed ledger technologies (DLTs).⁵⁴

The overarching goal of a DLT is to maintain a list of transactions that is agreed by all parties to be correct. This agreement is established via a consensus mechanism that entities using the DLT trust to provide a secure ledger. The choice of consensus mechanism depends on the

choice of DLT type, as well as its purpose. In the case of a permissionless public blockchain, a distributed ledger that anyone with an internet connection can read and write to, the consensus mechanism must eliminate any dependence on trust in fellow users to maintain a secure ledger. For this reason, many blockchains of this type must incentivise computational work by offering cryptocurrency in exchange for the generation of a cryptographically linked chain of transaction blocks (hence the name 'blockchain' for DLTs organised in this manner). This consensus mechanism, referred to as 'proof of work,' is effective at maintaining a trusted (but decentralised) ledger as long as no single entity, or group of cooperating entities, control the majority of the resources doing the computational work (Li, Jiang, Chen, Luo, & Wen, 2020).

Permissioned DLTs, on the other hand, restrict network access to authorised entities. In this case, the requirement to eliminate trust may be relaxed somewhat, and consensus mechanisms adapted appropriately—the transparency, immutability, and decentralised nature of the ledger is not dependent on a particular consensus mechanism. For example, in the case of strictly private DLTs that are used for record-keeping within an organisation, this mechanism may be as simple as having trusted members authorise any additions to the ledger. This form of DLT may be useful for large organisations who wish to keep all transactions of a particular type on one ledger that cannot be altered once recorded.

A particularly relevant type of permissioned DLT for the current work is federated blockchain, also known as a blockchain consortium. In its most general form, a blockchain consortium is a network of entities with a common interest in maintaining a shared record of transactions. Participating entities agree on a consensus protocol that will validate the transaction blocks, and access rights can be very flexible. While participating entities may add to this ledger (and also potentially validate transactions), there is no technical reason why the information on the ledger cannot be public. Note also that in the case of permissioned DLTs, it is often the case that entities on the network have a shared goal or use-case and, therefore, the incentive to participate in the consensus mechanism is simply the utility that each entity is able to extract through its access to the ledger—no cryptocurrency need be involved in the system.

At this point, it is necessary to directly address concerns that the types of DLTs that may be usefully implemented in patent systems may require excessive amounts of computational power. After all, many traditional blockchains, such as the Bitcoin blockchain, require a considerable amount of computational work to validate blocks of transactions (Krause & Tolaymat, 2018). This shortcoming is inherent to a proof-of-work based consensus mechanism as it is precisely this computational work that ensures that participants agree on the contents of the ledger (Nakamoto et al., 2008); however, this problem can be addressed with alternative consensus methods such as proof-of-stake and proof-of-authority in the case of permissionless DLTs (Wang et al., 2019). Permissioned DLTs, on the other hand, offer inherently computationally efficient validation and consensus mechanisms that are entirely different to those required for permissionless DLTs. This efficiency is possible because the identities of the transaction-validating nodes are known and they are generally few (Gramoli, 2020). Therefore, while complexity and computational inefficiency are generally significantly higher for *permissionless* DLTs than for traditional systems, we consider these particular downsides to be negligible when considering the implementation of a permissioned DLT.

Appendix B. The disclosure function of patents

The 'prospect theory' of patents (Kitch, 1977) highlights the many advantages of patents as claims of rights to exploit a particular piece of the technological landscape. Importantly, from the viewpoint of social welfare, significant efficiency gains in technological search are associated with disclosure (when compared to trade secrets): a patent signals the location of a technologically valuable piece of information, just as a mineral claim signals the geographical location of valuable minerals. In

⁵³ Within the context of this work, a 'transaction' generally refers to an interaction between an applicant or patentee and the patent office or examiner.

⁵⁴ DLTs, and blockchain in particular, are inextricably linked to cryptocurrency in the mind of the public. It is therefore worth noting that not one of the proposals suggested in this work is dependent on cryptocurrency. There is no technical reason why standard payment systems, using fiat currencies, could not be integrated into a DLT-based patent records system.

both scenarios, competitors will start searching, or prospecting, the surrounding landscape for a claim of their own while avoiding much duplicative effort searching less mineral-rich areas.⁵⁵ To optimise the societal-welfare-increasing function of patents, therefore, an ideal system would maximise the value of the disclosure mechanism without reducing the incentive for firms to partake in risky inventive activities—there is a need for balance between disclosure and incentives to innovate.

We can reframe disclosure requirements as existing on two axes: clarity and time. Clarity broadly reflects the usefulness of the information disclosed in the patent document—disclosure is not useful if information relevant to follow-on innovators is omitted, nor if it is obscured in legal terms unfamiliar to a ‘person having ordinary skill in the art’ (PHOSITA). Time reflects the length of the secrecy period that is included in the ‘grand bargain’—firms enjoy secrecy of their patented inventions for much of the inventions’ pendency. Even if no formal litigation takes place, this could result in years of expensive duplicative effort by competitors being, at least partially, wasted. At the extremes of these axes, conditional on the necessity of *some* form of disclosure consistent with the terms of the grand bargain, we find two scenarios: one that is optimal for the patenting firms, and one that is optimal for competitors and follow-on innovators. The former is when secrecy time is maximised (public at grant) and the invention description is allowed to be as obfuscated as possible such that it is of minimal assistance to a PHOSITA who would like to replicate the invention. The latter scenario is one in which an invention is public from the filing date such that secrecy time (that is, wasteful duplication) is minimised, while clarity is maximised by applicants being held to a very high standard such that a PHOSITA has all the information required to precisely replicate the invention. The *socially* optimal amount of disclosure likely lies somewhere between these two extremes, as we would like to encourage the use of the knowledge contained in granted patent without discouraging patenting in the first place and undermining their use as incentives for risky R&D. It may be argued that current standards of clarity lie too close to the extreme that is beneficial for the patenting firm (Fromer, 2008; Ouellette, 2012). Adjusting the time of disclosure as suggested in Section 4.1.1, therefore, may help compensate for this imbalance.

Appendix C. Secrecy orders

In many jurisdictions, when particular inventions disclosed to the patent office are determined to be a threat to national security, they may be placed under forced secrecy (e.g., 35 U.S.C. § 17 (United States); Patents Act 1977, Section 22 (United Kingdom)). Arguments for and against this practice notwithstanding (Lee, 1997; Relyea, 2002), the checks for inventions of this nature necessarily take place under current laws in many jurisdictions before any publication of the patent application or granted document. Because of this, publication on filing may not be feasible and instead should take place after these checks to ensure compliance with current law. Of course, full publication on filing may not be appropriate even in the absence of secrecy orders—the patent office would also need to check the formatting of patent applications is correct and classify the subject matter before publication.⁵⁶ Therefore, there are good arguments to have a delay from filing to ‘immediate’

⁵⁵ In the context of patents, Yelderman (2019) observed that about one-third of all patent invalidations resulting from an inter-partes review cited at least one patent or application that was secret at the time of filing, representing significant duplicative efforts and financial waste. According to the authors, “there was no lawful way for the inventor to discover the fatal prior art reference, no matter how hard she might have searched” (Yelderman, 2019). As inter partes review generally occurs after grant, one may hypothesise that many more applications are rejected during examination for a similar reason.

⁵⁶ It is still possible, however, to automatically publish a hash for every application on filing.

publication; however, we argue that this delay should be minimised and does not undermine the arguments laid out above. Moreover, these procedures are very jurisdiction dependent and, therefore, their interaction with local law should be examined on a case-by-case basis in any case. (Of course, this applies equally to all proposals presented in this work).

Appendix D. Trade secrets and DLTs

Also of note is the increasingly easy modes of keeping a trade secret while an invention is in development—if patents are being applied for at a later stage of development, then firms would be more dependent on trade secrets than they are currently. Firms offering blockchain-based trade secret records with linked non-disclosure agreements are already numerous⁵⁷ and can provide much more security and oversight than traditional systems. These trade secrets are well-defined and time-stamped, with recorded version histories and access logs, and thus can provide much more concrete evidence for misappropriation than in the past. Courts in China, for example, have already ruled that these documents are permissible as evidence in trade secret cases,⁵⁸ and the state legislature of Vermont in the United States has passed a law allowing blockchain records as evidence in state courts,⁵⁹ which presumably covers misappropriation of blockchain-based trade secrets. In the United States, third-party services are very relevant in light of recent changes to trade secret law included in the America Invents Act (112th Congress of the United States of America, 2011). If an entity makes use of a particular trade secret, and this trade secret is discovered independently and patented by a different entity, then the first entity is granted ‘prior usage rights’—provided they are able to prove that the patent invention was being used for at least a year before the filing of the patent. Proving the existence and use of this trade secret is made much simpler with DLTs as each record is time-stamped and immutable. Earlier disclosure is now more feasible than ever, thanks to the conjunction of these legal and technical developments.⁶⁰

Appendix E. Early filing and patent obfuscation

The ‘file early, file often’ attitude that exists in many industries leads inevitably to vague invention specifications. One solution to this problem, proposed by Cotropia (2009), is to add a patentability requirement stipulating that applicants reduce their invention to practice *before* examination (but not necessarily before application)—a kind of ‘brute force’ solution. This is certainly desirable, but inherently difficult to implement as it would add considerable administrative burden and subjectivity to the examination process in many cases.⁶¹ It would additionally disadvantage firms without the resources to fund the development of a functioning prototype before the investment that a

⁵⁷ Examples include Bernstein (bernstein.io) and iDefendo (idefendo.com) (all last accessed 10/12/2020).

⁵⁸ <https://tinyurl.com/blockchainChina> (last accessed 10/12/2020)

⁵⁹ <https://legislature.vermont.gov/statutes/section/12/081/01913> (last accessed 10/12/2020)

⁶⁰ Taken alone, this does not imply that the patent system itself requires a DLT component. However, it does suggest that the legal environment in which the patent system operates will likely become much more DLT-friendly to accommodate these changes.

⁶¹ Enablement becomes a particularly interesting requirement in this scenario as the lack of a proper invention specification enabling a PHOSITA to make and use the invention (without costly experimentation) becomes less acceptable if the invention has already been reduced to practice. Such problems may be solved by a kind of peer review system (Ouellette, 2012, 2016; Seymore, 2008), which is out of scope of the current work but may also be facilitated by DLT-based records and feedback (Ahn, Chang, Kokku, & Watson, 2019; Bartling, 2019).

patent may attract (Farre-Mensa et al., 2020), or those that intend to license the technology to firms with development capabilities as might universities (Jensen & Thursby, 2001; Owen-Smith & Powell, 2003).

To combat intentionally obscured patent specifications,⁶² there need to be higher standards for disclosure, particularly concerning enablement and the level of detail in the description of the technical principles on which the invention functions (Lemley, 2016). Claimed attributes and applications of inventions that exist in the real world (and have been tested for particular use cases) are, by definition, able to be specified more clearly and in more detail; resulting in narrower claims. In order to ensure the inventors that patent later in the development process are not disadvantaged in this way, strict rules must be placed on those that cannot produce evidence that the invention has been reduced to practice, such as limiting granted patents to the specified use cases (including prophetic examples), provided these are themselves not overly broad (Lemley, 2016).

It is additionally worth noting that early filing not only encumbers the patent office with dubious applications that amplify examination problems such as backlog and invalid grants.⁶³ In cases where a high-quality patent is granted and the invention commercialisable, the *patentee* now has much less time to recoup their investment in the development of the resultant product as the patent term starts at the filing date in most jurisdictions. Relatedly, filing earlier in the development process means that there is less time to fully develop the invention before it becomes financially unviable (due to a looming patent expiry date), resulting in ‘under-developed’ patent rights to the detriment of both the patentee and society, the latter of whom miss out on new-and-improved products (Abramowicz, 2006).

Appendix F. Office-specific considerations for prior art bounties

All proposals contained in this work are susceptible to office-specific limitations. Prior art bounties are a relatively simple suggestion with many components that are already implemented in some way in many patent offices. For these reasons, this proposal presents a good opportunity to illustrate how jurisdiction-specific procedures may interact with such a policy change.

The simple example involves a situation where two or more bounty-hunter submissions are used *in combination* to reject a patent application. At the EPO, there exists a ‘Y’ tag to indicate these kinds of citation. This event could require a split of the reward between bounty hunters when a combination contains multiple bounty-hunter submissions. The USPTO has a similar practice but does not have a formal tag to describe this scenario, further complicating matters. In any case, the way these cases are dealt with could change the incentives of the bounty hunters, particularly when submissions can only contain one piece of prior art at a time.

A more complex example concerns the ability for an applicant to amend an application before it is examined in response to a third-party prior art submission. This scenario would significantly reduce the incentive for bounty hunters to submit relevant prior art, as relevant art could be made irrelevant by an amendment and the bounty hunter would have wasted their time. We must also rule out the possibility of paying out the bounty should the applicant amend the document and submit the hunter’s prior art themselves as justification; this policy would constitute the use of the bounty hunter as a free search service paid for by the patent office. Ideally, we would like to incentivise firms to conduct their own in-depth prior art search *before* filing for a patent

that takes this art into account. For this reason, jurisdictions where applications can be amended before the search report or first office action may not extract as much utility as possible from prior art bounties.⁶⁴ In the case where amendment is not possible between the time the document is made public and the first office action, firms can still see suggested prior art, and prepare an amended document *in advance* to be submitted should the examiner use this information to justify a rejection; this could significantly reduce pendency times.

Assuming amendments are possible, the following may be an appropriate solution. If the applicant abandons or amends their application after the document is made public but before examination (in response to a prior art submission), the bounty-hunters’ submission fees are refunded. In the case where an application is withdrawn, then a small reward should also be given to hunters in acknowledgement of the resources they may have saved the patent office—as application fees are not refunded, and minimal resources have been spent on applications that are withdrawn before examination, the funds exist to provide such a reward.

Appendix G. Proof of existence of prior art submissions

From a technical point of view, any third party submitting relevant non-patent literature as prior art to the patent office in the capacity of a bounty hunter would not generally be able to *prove* that the art was publicly available before the patent application. This deficiency exists for current validity challenges and is not remedied with the DLT-based patent system proposed in Section 4.2.2. However, this is a need that would be catered to by third-party services via DLT-based defensive publications, working papers, preprints, or other technical documents—these services already exist in forms that are low-cost and easily accessible.⁶⁵

Appendix H. Gold-plated patents, SPER, and private examiner incentives

One idea to reduce levels of unnecessary patent litigation, popularised by Mark Lemley and colleagues (Lemley, 2011; Lemley et al., 2005; Lichtman & Lemley, 2007) and directed at the USPTO, proposes a weakening of the default presumption of validity and encourages a move towards an *earned* presumption of validity.⁶⁶ This is sometimes framed as the ‘gold-plating’ of a patent (Lemley et al., 2005; Lichtman & Lemley, 2007), and would simply grant a presumption of validity to those patents that survive additional scrutiny. This review process would be conducted at the request of, and paid for by, either an applicant desiring additional certainty in their patent rights or by a competitor who believes the patent application should be rejected under the patent office’s own standards (similarly to current inter-partes and post-grant reviews). This mechanism ensures that only particularly valuable patents undergo this additional scrutiny.

The above proposal dovetails nicely with another popular proposal: ‘second-pair-of-eyes’ review (SPER). This process, which would require two examiners to agree on the patentability of a particular application, was actually implemented at the USPTO in the year 2000 for a particular class of business methods patents; being somewhat controversial as

⁶² This behaviour likely occurs within the current system as well (Feldman, 2008; McJohn, 2007).

⁶³ If new developments lie outside the scope of the original application, then follow-up applications such as continuations may be filed that all but supersede the initial filing, adding to the already large backlog of patent applications (Cotropia et al., 2013).

⁶⁴ At the EPO, applications may be amended between receipt of the search report and before examination, while at the USPTO, applications may be amended at any point before the first office action (and under certain conditions after this point).

⁶⁵ See, e.g., blockchainyourip.com, bernstein.io, artifacts.ai, and idefendo.com (all last accessed 10/12/2020).

⁶⁶ U.S. courts automatically presume that the USPTO issues valid patents, which means that it is the responsibility of those challenging the patent to provide clear and convincing evidence that the patent is *invalid* (Seymore, 2012).

patentable matter in the first place (Hall, 2003; Spulber, 2011), and attracting many applications of questionable quality, this class appeared to be a suitable test-bed for SPER (Allison & Hunter, 2006). This pilot resulted in a dramatic drop in grant rates for this class (down to 17%, compared to about 70% for most other classes) according to Lemley and Sampat (2008), indicating that the scheme may have indeed filtered out a high number of unpatentable inventions. At the same time, however, the high technological similarities between patent classifications meant that it was generally not difficult to draft applications in a way that avoided a main classification⁶⁷ that would result in a SPER (Allison & Hunter, 2006).

Meanwhile, a granted patent may be challenged by third parties, or even the applicant themselves if they feel the resulting patent has been unnecessarily narrowed. These challenges (and indeed appeals, in the case of rejected applications) are generally handled in the United States by the Patent Trial and Appeal Board and in Europe by the Opposition Divisions of the EPO (Marsnik, 2013). While this process would remain an important check in the presence of third-party examiners, the patent office is currently not held accountable for granting invalid patents (Thomas, 2002), or indeed rejecting valid applications. Therefore, with a contractor monitoring system in place and the patent office having the power to terminate the certification of these contractors, third party examiners already have a much larger incentive to make correct patentability decisions than do current examiners. After all, the risk the contractor takes on when accepting a contract is that their decisions are found to be incorrect on review⁶⁸—this factor is almost entirely in their control and so is not subject to the usual risks associated with government contracts such as higher than expected costs or low ex-post demand for the end-product (Jensen & Stonecash, 2005; Laffont & Tirole, 1993). In fact, it may be argued that the patent office taking on the responsibility for penalising contractors for poor performance constitutes a much higher level of accountability than that at present.

References

- 112th Congress of the United States of America. (2011). Pub. L. no. 112-29, 125 stat. 284. <https://www.govinfo.gov/app/details/PLAW-112publ29>.
- 35 U.S.C. 17 (United States). (2006). Invention Secrecy Act of 1951. <https://www.govinfo.gov/app/details/USCODE-2011-title35/USCODE-2011-title35-partII-chap17>.
- Abramowicz, M. (2006). The danger of underdeveloped patent prospects. *Cornell Law Review*, 92, 1065.
- Abramowicz, M., & Duffy, J. F. (2008). Ending the patenting monopoly. *University of Pennsylvania Law Review*, 157, 1541.
- Ackerman, L. J. (2011). Prioritization: Addressing the patent application backlog at the United States patent and trademark office. *Berkeley Technology Law Journal*, 26(1), 67–92.
- Ahn, J.-W., Chang, M. D., Kokku, R., & Watson, P. (June 11 2019). *Blockchain for open scientific research*. US Patent App. 10/320,574.
- Allen, D. W. E., Berg, C., Davidson, S., & Potts, J. (2020). *Property rights, knowledge commons and blockchain governance*. Available at SSRN 3344110.
- Allison, J. R., & Hunter, S. D. (2006). On the feasibility of improving patent quality one technology at a time: The case of business methods. *Berkeley Technology Law Journal*, 21, 729.
- Allison, J. R., & Lemley, M. A. (1998). Empirical evidence on the validity of litigated patents. *AIPLA Quarterly Journal*, 26, 185.
- Anderson, J. J. (2016). Nontechnical disclosure. *Vanderbilt Law Review*, 69, 1573.
- Anderson, N. P. (2015). Striking a balance: The pursuit of transparent patent ownership. *Berkeley Technology Law Journal*, 30(4), 395–444.
- Androulaki, E., Barger, A., Bortnikov, V., Cachin, C., Christidis, K., De Caro, A., ... Manevich, Y., et al. (2018). Hyperledger fabric: A distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference* (p. 30). ACM.
- Aoki, R., & Spiegel, Y. (2009). Pre-grant patent publication and cumulative innovation. *International Journal of Industrial Organization*, 27(3), 333–345.
- Arora, A., Fosfuri, A., & Gambardella, A. (2001). Markets for technology and their implications for corporate strategy. *Industrial and Corporate Change*, 10(2), 419–451.
- Arora, A., & Gambardella, A. (2010). Ideas for rent: An overview of markets for technology. *Industrial and Corporate Change*, 19(3), 775–803.
- Arundel, A. (2001). The relative effectiveness of patents and secrecy for appropriation. *Research Policy*, 30(4), 611–624.
- Baliga, A. (2017). Understanding blockchain consensus models. In *Persistent systems*.
- Bartling, S. (2019). Blockchain for science and knowledge creation. In *Gesundheit digital* (pp. 159–180). Springer.
- Baudry, M., & Dumont, B. (2006). Patent renewals as options: Improving the mechanism for weeding out lousy patents. *Review of Industrial Organization*, 28(1), 41–62.
- Benčić, F. M., & Žarko, I. P. (2018). Distributed ledger technology: Blockchain compared to directed acyclic graph. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)* (pp. 1569–1570). IEEE.
- Benet, J. (2014). *IPFS - content addressed, versioned, P2P file system*. arXiv preprint arXiv:1407.3561.
- Berg, C., Davidson, S., & Potts, J. (2019). *Understanding the blockchain economy: An introduction to institutional cryptoeconomics*. Edward Elgar Publishing.
- Bessen, J., Ford, J., & Meurer, M. J. (2011). The private and social costs of patent trolls. *Regulation*, 34, 26.
- Bessen, J., & Meurer, M. J. (2008). *Patent failure: How judges, bureaucrats, and lawyers put innovators at risk*. Princeton University Press.
- Bessen, J., & Meurer, M. J. (2013). The patent litigation explosion. *Loyola University Chicago Law Journal*, 45, 401.
- Blind, K., Cremers, K., & Mueller, E. (2009). The influence of strategic patenting on companies' patent portfolios. *Research Policy*, 38(2), 428–436.
- Blind, K., Edler, J., Frietsch, R., & Schmoch, U. (2006). Motives to patent: Empirical evidence from Germany. *Research Policy*, 35(5), 655–672.
- Boldrin, M., & Levine, D. K. (2008). *Against intellectual monopoly*. Cambridge: Cambridge University Press.
- Boldrin, M., & Levine, D. K. (2013). The case against patents. *Journal of Economic Perspectives*, 27(1), 3–22.
- Brouwer, E., & Kleinknecht, A. (1999). Innovative output, and a firm's propensity to patent: An exploration of CIS micro data. *Research Policy*, 28(6), 615–624.
- Buenstorf, G., & Geissler, M. (2012). Not invented here: Technology licensing, knowledge transfer and innovation based on public research. *Journal of Evolutionary Economics*, 22(3).
- Cabral, J. E., Chavan, A., Clarke, T. M., & Greacen, J. (2012). Using technology to enhance access to justice. *Harvard Journal of Law & Technology*, 26, 241.
- Cachin, C., & Vukolić, M. (2017). *Blockchain consensus protocols in the wild*. arXiv preprint arXiv:1707.01873.
- Capponi, G., Criscuolo, P., Martinelli, A., & Nuvolari, A. (2019). Profiting from innovation: Evidence from a survey of Queen's awards winners. *Structural Change and Economic Dynamics*, 49, 155–169.
- Catalini, C., & Gans, J. S. (2016). Some simple economics of the blockchain. In *National Bureau of Economic Research*, (w22952).
- Chien, C. V. (2010). From arms race to marketplace: The complex patent ecosystem and its implications for the patent system. *Hastings Law Journal*, 62, 297.
- Choi, J. P. (2005). Live and let live: A tale of weak patents. *Journal of the European Economic Association*, 3(2–3), 724–733.
- Cohen, W. M., Goto, A., Nagata, A., Nelson, R. R., & Walsh, J. P. (2002). R&D spillovers, patents and the incentives to innovate in Japan and the United States. *Research Policy*, 31(8–9), 1349–1367.
- Cohen, W. M., Nelson, R. R., & Walsh, J. P. (2000). Protecting their intellectual assets: Appropriability conditions and why us manufacturing firms patent (or not). In *National Bureau of Economic Research*, (w7552).
- Conti, A., Thursby, J., & Thursby, M. (2013). Patents as signals for startup financing. *The Journal of Industrial Economics*, 61(3), 592–622.
- Cornelli, F., & Schankerman, M. (1999). Patent renewals and R&D incentives. *RAND Journal of Economics*, 197–213.
- Cotropia, C. A. (2009). The folly of early filing in patent law. *Hastings Law Journal*, 61 (65).
- Cotropia, C. A., Quillen, C. D., Jr., & Webster, O. H. (2013). Patent applications and the performance of the US patent and trademark office. *Federal Circuit Bar Journal*, 23, 179.
- Czarnitzki, D., Hussinger, K., & Schneider, C. (2015). R&D collaboration with uncertain intellectual property rights. *Review of Industrial Organization*, 46(2), 183–204.
- Dai, J., & Vasarhelyi, M. A. (2017). Toward blockchain-based accounting and assurance. *Journal of Information Systems*, 31(3), 5–21.
- Domberger, S., & Jensen, P. (1997). Contracting out by the public sector: Theory, evidence, prospects. *Oxford Review of Economic Policy*, 13(4), 67–78.
- Draho, P. (2010). *The global governance of knowledge: Patent offices and their clients*. Cambridge University Press.
- Dunphy, P., & Petitcolas, F. A. (2018). A first look at identity management schemes on the blockchain. *IEEE Security and Privacy*, 16(4), 20–29.
- Eckert, A., & Langinier, C. (2014). A survey of the economics of patent systems and procedures. *Journal of Economic Surveys*, 28(5), 996–1015.
- El Ioini, N., & Pahl, C. (2018). A review of distributed ledger technologies. In *OTM Confederated International Conferences "On the Move to Meaningful Internet Systems"* (pp. 277–288). Springer.
- Farrell, J., & Shapiro, C. (2008). How strong are weak patents? *American Economic Review*, 98(4), 1347–1369.
- Farre-Mensa, J., Hegde, D., & Ljungqvist, A. (2020). What is a patent worth? Evidence from the U.S. patent "lottery". *The Journal of Finance*, 75(2), 639–682.
- Feldman, R. (2008). Plain language patents. *Texas Intellectual Property Law Journal*, 17, 289.
- Frakes, M. D., & Wasserman, M. F. (2015). Does the US patent and trademark office grant too many bad patents: Evidence from a quasi-experiment. *Stanford Law Review*, 67, 613.

⁶⁷ Applications with the SPER-enforced class as a secondary classification were exempt.

⁶⁸ Both incorrect rejections and incorrect grants.

- Frakes, M. D., & Wasserman, M. F. (2017). Is the time allocated to review patent applications inducing examiners to grant invalid patents? Evidence from microlevel application data. *Review of Economics and Statistics*, 99(3), 550–563.
- Frakes, M. D., & Wasserman, M. F. (2019). Irrational ignorance at the patent office. *Vanderbilt Law Review*, 72.
- Freilich, J. (2019). Prophetic patents. *UCDL Rev.*, 53, 663.
- Freilich, J., & Ouellette, L. L. (2019). Science fiction: Fictitious experiments in patents. *Science*, 364(6445), 1036–1037.
- Fridgen, G., Radszuwili, S., Urbach, N., & Utz, L. (2018). Cross-organizational workflow management using blockchain technology-towards applicability, auditability, and automation. In *Proceedings of the 51st Hawaii International Conference on System Sciences*.
- Fromer, J. C. (2008). Patent disclosure. *Iowa Law Review*, 94, 539.
- Fromer, J. C. (2016). Dynamic patent disclosure. *Vanderbilt Law Review*, 69, 1715.
- Furman, J. L., Nagler, M., & Watzinger, M. (2018). Disclosure and subsequent innovation: Evidence from the patent depository library program. In *National Bureau of Economic Research*, (w24660).
- Gallini, N. T., & Wright, B. D. (1990). Technology transfer under asymmetric information. *The Rand Journal of Economics*, 147–160.
- Gambardella, A., Giuri, P., & Luzzi, A. (2007). The market for patents in Europe. *Research Policy*, 36(8), 1163–1183.
- Gambardella, A., Harhoff, D., & Verspagen, B. (2008). The value of European patents. *European Management Review*, 5(2), 69–84.
- Gans, J. S., Hsu, D. H., & Stern, S. (2008). The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays. *Management Science*, 54(5), 982–997.
- Gans, J. S., King, S. P., & Lampe, R. (2004). Patent renewal fees and self-funding patent offices. *Topics in Theoretical Economics*, 4(1).
- GAO (United States Government Accountability Office). (2016). Patent office should strengthen search capabilities and better monitor examiners' work. In *GAO-16-479*.
- Ghafele, R., Gilbert, B., & DiGiammarino, P. (2011). Driving innovation through patent application review: The power of crowdsourcing prior art search. *Journal of Intellectual Property Rights*, 303–308.
- Gong, Y., Yang, J., & Shi, X. (2020). Towards a comprehensive understanding of digital transformation in government: Analysis of flexibility and enterprise architecture. *Government Information Quarterly*, 37(3), Article 101487.
- Graham, S., & Hegde, D. (2015). Disclosing patents' secrets. *Science*, 347(6219), 236–237.
- Graham, S. J., Marco, A. C., & Myers, A. F. (2018). Patent transactions in the marketplace: Lessons from the USPTO patent assignment dataset. *Journal of Economics and Management Strategy*, 27(3), 343–371.
- Gramoli, V. (2020). From blockchain consensus back to byzantine consensus. *Future Generation Computer Systems*, 107, 760–769.
- Hall, B. H. (2003). Business method patents, innovation, and policy. In *National Bureau of economic research*, (w9717).
- Hall, B. H. (2004). Exploring the patent explosion. *Journal of Technology Transfer*, 30 (1–2), 35–48.
- Hall, B. H. (2019). Is there a role for patents in the financing of new innovative firms? *Industrial and Corporate Change*, 28(3), 657–680.
- Hall, B. H., & Harhoff, D. (2012). Recent research on the economics of patents. *Annual Review of Economics*, 4(1), 541–565.
- Hall, B. H., Helmers, C., & von Graevenitz, G. (2015). Technology entry in the presence of patent thickets. In *National Bureau of Economic Research WP No 21455*. NBER.
- Hall, B. H., & Ziedonis, R. H. (2001). The patent paradox revisited: An empirical study of patenting in the US semiconductor industry, 1979–1995. *RAND Journal of Economics*, 101–128.
- Harabi, N. (1995). Appropriability of technical innovations an empirical analysis. *Research Policy*, 24(6), 981–992.
- Häussler, C., Harhoff, D., & Müller, E. (2012). To be financed or not: The role of patents for venture capital-financing. In *ZEW-Centre for European Economic Research Discussion Paper*, (09–003).
- Hegde, D., & Luo, H. (2017). Patent publication and the market for ideas. *Management Science*, 64(2), 652–672.
- Heger, D., & Hussinger, K. (2017). Implications of uncertain patent rights for german start-ups' commercialisation activities and access to external capital. *Industry and Innovation*, 24(7), 753–773.
- Hemel, D. J., & Ouellette, L. L. (2019). Innovation policy pluralism. *The Yale Law Journal*, 128, 544.
- Henkel, J., & Zischka, H. (2019). How many patents are truly valid? Extent, causes, and remedies for latent patent invalidity. *European Journal of Law and Economics*, 48(2), 195–239.
- Holden, R. T., & Malani, A. (2019). Can blockchain solve the hold-up problem in contracts?. In *National Bureau of Economic Research*, (w25833).
- Hsu, D. H., & Ziedonis, R. H. (2008). Patents as quality signals for entrepreneurial ventures. In *Academy of Management Proceedings*, number 1 in 2008 (pp. 1–6). Academy of Management.
- Hurwicz, L. (1973). The design of mechanisms for resource allocation. *American Economic Review*, 63(2), 1–30.
- IP Australia. (2018). Patent manual of practice & procedure. http://manuals.ipaustralia.gov.au/patents/national/examination/2.13.11.1_Notifications.htm. Retrieved 30/08/2019.
- Jaffe, A. B., & Lerner, J. (2004). *Innovation and its discontents: How our broken patent system is endangering innovation and Progress, and what to do about it*. Princeton University Press.
- Janssen, M., Matheus, R., Longo, J., & Weerakkody, V. (2017). Transparency-by-design as a foundation for open government. *Transforming Government: People, Process and Policy*, 11(1), 2–8.
- Jeffery, J. A. (2002). Preserving the presumption of patent validity: An alternative to outsourcing the US patent examiner's prior art search. *Catholic University Law Review*, 52, 761.
- Jensen, P. H., Palangkaraya, A., & Webster, E. (2015). Trust and the market for technology. *Research Policy*, 44(2), 340–356.
- Jensen, P. H., & Stonecash, R. E. (2005). Incentives and the efficiency of public sector-outsourcing contracts. *Journal of Economic Surveys*, 19(5), 767–787.
- Jensen, R., & Thursby, M. (2001). Proofs and prototypes for sale: The licensing of university inventions. *American Economic Review*, 91(1), 240–259.
- Jin, X., Spangler, S., Chen, Y., Cai, K., Ma, R., Zhang, L., ... Han, J. (2011). Patent maintenance recommendation with patent information network model. In *2011 IEEE 11th International Conference on Data Mining* (pp. 280–289). IEEE.
- Kannengießer, N., Lins, S., Dehling, T., & Sunyaev, A. (2020). Trade-offs between distributed ledger technology characteristics. *ACM Computing Surveys (CSUR)*, 53(2), 1–37.
- Kazenske, E. R. (2003). The future of prior art searching at the United States patent and trademark office. *World Patent Information*, 25(4), 283–287.
- Kitch, E. W. (1977). The nature and function of the patent system. *The Journal of Law and Economics*, 20(2), 265–290.
- Krause, M. J., & Tolaymat, T. (2018). Quantification of energy and carbon costs for mining cryptocurrencies. *Nature Sustainability*, 1(11), 711.
- Kshetri, N. (2018). Blockchain's roles in meeting key supply chain management objectives. *International Journal of Information Management*, 39, 80–89.
- Kuperberg, M. (2020). Blockchain-based identity management: A survey from the enterprise and ecosystem perspective. *IEEE Transactions on Engineering Management*, 67(4), 1008–1027.
- Laffont, J.-J., & Tirole, J. (1993). *A theory of incentives in procurement and regulation*. MIT Press.
- Lee, S. H. (1997). Protecting the private inventor under the peacetime provisions of the invention secrecy act. *Berkeley Technology Law Journal*, 12, 345.
- Lemley, M. A. (2000). Reconceiving patents in the age of venture capital. *Journal of Small & Emerging Business Law*, 4, 137.
- Lemley, M. A. (2011). Can the patent office be fixed. *Marquette Intellectual Property Law Review*, 15, 295.
- Lemley, M. A. (2016). Ready for patenting. *Boston University Law Review*, 96, 1171.
- Lemley, M. A., Lichtman, D., & Sampat, B. N. (2005). What to do about bad patents. *Regulation*, 28(4), 10–13.
- Lemley, M. A., & Myhrvold, N. (2007). How to make a patent market. *Hofstra Law Review*, 36, 257.
- Lemley, M. A., Richardson, K., & Oliver, E. (2018). The patent enforcement iceberg. *Texas Law Review*, 97, 801.
- Lemley, M. A., & Sampat, B. (2008). Is the patent office a rubber stamp. *The Emory Law Journal*, 58, 181.
- Lemley, M. A., & Shapiro, C. (2005). Probabilistic patents. *Journal of Economic Perspectives*, 19(2), 75–98.
- Lerner, J., & Tirole, J. (2015). Standard-essential patents. *Journal of Political Economy*, 123(3), 547–586.
- Lesavre, L., Varin, P., Mell, P., Davidson, M., & Shook, J. (2019). A taxonomic approach to understanding emerging blockchain identity management systems. *arXiv preprint arXiv:1908.00929*.
- Levin, R. C., Klevorick, A. K., Nelson, R. R., Winter, S. G., Gilbert, R., & Griliches, Z. (1987). Appropriating the returns from industrial research and development. *Brookings Papers on Economic Activity*, 1987(3), 783–831.
- Li, X., Jiang, P., Chen, T., Luo, X., & Wen, Q. (2020). A survey on the security of blockchain systems. *Future Generation Computer Systems*, 107, 841–853.
- Lichtman, D., & Lemley, M. A. (2007). Rethinking patent law's presumption of validity. *Stanford Law Review*, 60(45).
- Long, C. (2002). Patent signals. *The University of Chicago Law Review*, 625–679.
- Long, P. O. (1991). Invention, authorship, "intellectual property," and the origin of patents: Notes toward a conceptual history. *Technology and Culture*, 32(4), 846–884.
- Lück, S., Balsmeier, B., Seliger, F., & Fleming, L. (2020). Early disclosure of invention and reduced duplication: An empirical test. *Management Science*, 66(6), 2677–2685.
- Macdonald, S. (2004). When means become ends: Considering the impact of patent strategy on innovation. *Information Economics and Policy*, 16(1), 135–158.
- Marschall, S. (2020). Concept for a multinational patent search using elements of blockchain technology. *GRUR International*, 69(3), 225–232.
- Marsnik, S. J. (2013). Will the America invents act post-grant review improve the quality of patents? A comparison with the European patent office opposition. In *The changing face of US patent law and its impact on business strategy* (p. 183).
- McJohn, S. M. (2007). Patents: Hiding from history. *Santa Clara Computer & High Technology Law Journal*, 24, 961.
- Mejer, M., & van Pottelsberghe, B. (2011). Patent backlogs at uspto and epo: Systemic failure vs deliberate delays. *World Patent Information*, 33(2), 122–127.
- Menell, P. S., & Meurer, M. J. (2013). Notice failure and notice externalities. *Journal of Legal Analysis*, 5(1), 1–59.
- Mingxiao, D., Xiaofeng, M., Zhe, Z., Xiangwei, W., & Qijun, C. (2017). A review on consensus algorithm of blockchain. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)* (pp. 2567–2572). IEEE.
- Mirowski, P. (2011). *Science-mart*. Harvard University Press.
- Moore, K. A. (2005). Worthless patents. *Berkeley Tech. LJ*, 20, 1521.
- Nakamoto, S., et al. (2008). *Bitcoin: A peer-to-peer electronic cash system*.
- Nam, T. (2012). Suggesting frameworks of citizen-sourcing via government 2.0. *Government Information Quarterly*, 29(1), 12–20.

- Noveck, B. S. (2006). Peer to patent: Collective intelligence, open review, and patent reform. *Harvard Journal of Law & Technology*, 20, 123.
- Ølnes, S., & Jansen, A. (2017). Blockchain technology as a support infrastructure in e-government. In *International Conference on Electronic Government* (pp. 215–227). Springer.
- Ølnes, S., Ubacht, J., & Janssen, M. (2017). Blockchain in government: Benefits and implications of distributed ledger technology for information sharing. *Government Information Quarterly*, 34(3), 355–364.
- Ouellette, L. L. (2012). Do patents disclose useful information? *Harvard Journal of Law & Technology*, 25(2), 545.
- Ouellette, L. L. (2016). Pierson, peer review, and patent law. *Vanderbilt Law Review*, 69, 1825.
- Owen-Smith, J., & Powell, W. W. (2003). The expanding role of university patenting in the life sciences: Assessing the importance of experience and connectivity. *Research Policy*, 32(9), 1695–1711.
- Pakes, A. (1986). Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica: Journal of the Econometric Society*, 755–784.
- Patents Act 1977 (United Kingdom). Patents Act 1977, c. 22. Available at: <https://tinyurl.com/UKPatentSecrecy>.
- Pervez, H., Muneeb, M., Irfan, M. U., & Haq, I. U. (2018). A comparative analysis of DAG-based blockchain architectures. In *2018 12th International Conference on Open Source Systems and Technologies (ICOSST)* (pp. 27–34). IEEE.
- Pitts, A., & Kim, J. (2009). The patent prosecution highway: Is life in the fast lane worth the cost. *Hastings Sci. & Tech. LJ*, 1, 127.
- Popov, S. (2016). *The tangle*. IOTA Whitepaper.
- Prager, F. D. (1944). A history of intellectual property from 1545 to 1787. *Journal of the Patent Office Society*, 26, 711.
- Rai, A. K. (2008). Growing pains in the administrative state: The patent office's troubled quest for managerial control. *University of Pennsylvania Law Review*, 157, 2051.
- de Rassenfosse, G. (2012). How SMEs exploit their intellectual property assets: Evidence from survey data. *Small Business Economics*, 39(2), 437–452.
- de Rassenfosse, G. (2018). Notice failure revisited: Evidence on the use of virtual patent marking. In *National Bureau of Economic Research*, (w24288).
- de Rassenfosse, G., Griffiths, B., Jaffe, A. B., & Webster, E. (2020). Low-quality patents in the eye of the beholder: Evidence from multiple examiners. *Journal of Law, Economics, and Organization* (forthcoming).
- de Rassenfosse, G., & Higham, K. (2020). Wanted: A standard for virtual patent marking. *Journal of Intellectual Property Law & Practice*, 15(7), 544–553.
- de Rassenfosse, G., & Jaffe, A. B. (2018). Are patent fees effective at weeding out low-quality patents? *Journal of Economics and Management Strategy*, 27(1), 134–148.
- de Rassenfosse, G., & Raiteri, E. (2020). Technology protectionism and the patent system: Strategic technologies in China. *Journal of Industrial Economics*. Forthcoming.
- de Rassenfosse, G., & van Pottelsberghe de la Potterie, B. (2013). The role of fees in patent systems: Theory and evidence. *Journal of Economic Surveys*, 27(4), 696–716.
- Rauchs, M., Glidden, A., Gordon, B., Pieters, G. C., Recanatini, M., Rostand, F., ... Zhang, B. Z. (2018). *Distributed ledger technology systems: A conceptual framework*. Available at SSRN 3230013.
- Relyea, H. C. (2002). Homeland security and information. *Government Information Quarterly*, 19(3), 213–223.
- Remus, D., & Levy, F. (2017). Can robots be lawyers: Computers, lawyers, and the practice of law. *Georgetown Journal of Legal Ethics*, 30, 501.
- Righi, C., & Simcoe, T. (2019). Patent examiner specialization. *Research Policy*, 48(1), 137–148.
- Saberi, S., Koushizadeh, M., Sarkis, J., & Shen, L. (2019). Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, 57(7), 2117–2135.
- Saidi, F., & Zaldokas, A. (2020). How does firms' innovation disclosure affect their banking relationships? In *Management Science*. Forthcoming.
- Sampat, B. (2005). Determinants of patent quality: An empirical analysis. Available at <https://www.immagick.com/eLibrary/ARCHIVES/GENERAL/COLUMBIA/C050902S.pdf>.
- Sampat, B. N. (2018). A survey of empirical evidence on patents and innovation. In *National Bureau of economic research*, (w25383).
- Savelyev, A. (2018). Copyright in the blockchain era: Promises and challenges. *Computer Law and Security Review*, 34(3), 550–561.
- Schuett, F. (2013). Patent quality and incentives at the patent office. *The Rand Journal of Economics*, 44(2), 313–336.
- Scotchmer, S. (1999). On the optimality of the patent renewal system. *The Rand Journal of Economics*, 181–196.
- Seymore, S. B. (2008). Heightened enablement in the unpredictable arts. *UCLA Law Review*, 56, 127.
- Seymore, S. B. (2012). The presumption of patentability. *Minnesota Law Review*, 97, 990.
- Shapiro, C. (2000). Navigating the patent thicket: Cross licenses, patent pools, and standard setting. *Innovation Policy and the Economy*, 1, 119–150.
- Shrier, D., Wu, W., & Pentland, A. (2016). Blockchain & infrastructure (identity, data security). *Massachusetts Institute of Technology-Connection Science*, 1(3), 1–19.
- Sichelman, T., & O'Connor, S. (2012). Patents as promoters of competition: The guild origins of patent law in the Venetian Republic. *San Diego Law Review*, 49, 1267.
- Sohn, J. L. (2011). Can't the PTO get a little respect. *Berkeley Technology Law Journal*, 26, 1603.
- Sousa, J., Bessani, A., & Vukolic, M. (2018). A byzantine fault-tolerant ordering service for the hyperledger fabric blockchain platform. In *2018 48th annual IEEE/IFIP international conference on dependable systems and networks (DSN)* (pp. 51–58). IEEE.
- Spulber, D. F. (2011). Should business method inventions be patentable? *Journal of Legal Analysis*, 3(1), 265–340.
- Stiglitz, J. E. (2007). Economic foundations of intellectual property rights. *Duke Law Journal*, 57, 1693.
- Strandburg, K. J., Csardi, G., Tobochnik, J., & Erdi, P. (2006). Law and the science of networks: An overview and an application to the patent explosion. *Berkeley Technology Law Journal*, 21, 1293.
- Swan, M. (2015). *Blockchain: Blueprint for a new economy*. O'Reilly Media, Inc.
- Tabakovic, H., & Wollmann, T. G. (2018). From revolving doors to regulatory capture? Evidence from patent examiners. In *National Bureau of Economic Research*, (w24638).
- Thomas, J. R. (2001). Collusion and collective action in the patent system: A proposal for patent bounties. In *Entrepreneurial Inputs and Outcomes: New studies of entrepreneurship in the United States* (pp. 95–132). Emerald Group Publishing Limited.
- Thomas, J. R. (2002). The responsibility of the rulemaker: Comparative approaches to patent administration reform. *Berkeley Technology Law Journal*, 17, 727.
- Thurrow, L. C. (1997). Needed: A new system of intellectual property rights. *Harvard Business Review*, 75, 94–107.
- Trimble, M. (2016). Patent working requirements: Historical and comparative perspectives. *UC Irvine Law Review*, 6, 483.
- United States Patent and Trademark Office. (2003). USPTO 21st century strategic plan. <https://www.uspto.gov/about-us/performance-and-planning/strategic-planning/st-rategic-plan-21st-century>.
- United States Senate. (1992). The Patent System Harmonization Act of 1992: Joint hearing before the Subcommittee on Patents, Copyrights, and Trademarks of the Senate Committee on the Judiciary. <https://catalog.hathitrust.org/Record/008517993>.
- USPTO. (2014). Report on virtual marking. https://www.uspto.gov/sites/default/files/aia_implementation/VMreport.pdf.
- Veer, T., & Jell, F. (2012). Contributing to markets for technology? A comparison of patent filing motives of individual inventors, small companies and universities. *Technovation*, 32(9–10), 513–522.
- Vonneuman, N. A. (1956). Conditionally exclusive patent rights and the patent clause of the constitution. *American Journal of Comparative Law*, 5, 391.
- Wagner, R. P. (2008). Understanding patent-quality mechanisms. *University of Pennsylvania Law Review*, 157, 2135.
- Wang, W., Hoang, D. T., Hu, P., Xiong, Z., Niyato, D., Wang, P., ... Kim, D. I. (2019). A survey on consensus mechanisms and mining strategy management in blockchain networks. *IEEE Access*, 7, 22328–22370.
- Wasserman, M. F. (2012). The changing guard of patent law: Chevron deference for the PTO. *William & Mary Law Review*, 54, 1959.
- Wijnhoven, F., Ehrenhard, M., & Kuhn, J. (2015). Open government objectives and participation motivations. *Government Information Quarterly*, 32(1), 30–42.
- WIPO. (2012). Administrative instructions under the patent cooperation treaty. https://www.wipo.int/pct/en/texts/ai/s801.html#_s801_a. Retrieved 30/08/2019.
- Wüst, K., & Gervais, A. (2018). Do you need a blockchain? In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)* (pp. 45–54). IEEE.
- Xia, Q., Sifah, E. B., Asamoah, K. O., Gao, J., Du, X., & Guizani, M. (2017). Medshare: Trust-less medical data sharing among cloud service providers via blockchain. *IEEE Access*, 5, 14757–14767.
- Xu, X., Weber, I., Staples, M., Zhu, L., Bosch, J., Bass, L., ... Rimba, P. (2017). A taxonomy of blockchain-based systems for architecture design. In *2017 IEEE International Conference on Software Architecture* (pp. 243–252). IEEE.
- Yamauchi, I., & Nagaoka, S. (2015). Does the outsourcing of prior art search increase the efficiency of patent examination? Evidence from Japan. *Research Policy*, 44(8), 1601–1614.
- Yang, W., Garg, S., Raza, A., Herbert, D., & Kang, B. (2018). Blockchain: Trends and future. In *Pacific Rim Knowledge Acquisition Workshop* (pp. 201–210). Springer.
- Yelderman, S. (2019). Prior art in inter partes review. *Iowa Law Review*, 104, 2705–2734.
- Zavolokina, L., Ziolkowski, R., Bauer, I., & Schwabe, G. (2020). Management, governance and value creation in a blockchain consortium. *MIS Quarterly Executive*, 19(1), 1–17.