Decentralising the Patent System

Gaétan de Rassenfosse¹ and Kyle Higham¹

¹Ecole polytechnique fédérale de Lausanne, Switzerland

December 17, 2019

Abstract

This paper proposes a substantive re-think of the modern patent system. The patent system has come under intensive criticism in the past, and many scholars have proposed ways to improve it. Ideas for improvement include, e.g., prior-art bounties, contracting out examination and dynamic fee setting. However, many of these ideas have gone unheeded due to the cost of administering them and the rigidity of the patent system. We explore how distributed ledger technologies enable these major changes.

Keywords— blockchain, distributed ledger, intellectual property, patent, smart contract

JEL: K11, K23, L24, O34, O38

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 (Prager, 1944; Long, 1991; Sichelman and 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 (Thurow, 1997; Jaffe and Lerner, 2004; Stiglitz, 2007; Bessen and Meurer, 2008), and additionally come with such high costs to society that many have wondered why governments should grant these rights at all (Boldrin and Levine, 2008, 2013; Mirowski, 2011).

In this work, we argue that distributed ledger technologies (DLTs) may improve the effectiveness of modern patent systems, and illustrate this point with vivid 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 (Gans et al., 2008; Ackerman, 2011), high levels of patent invalidity (Lemley and Shapiro, 2005; Frakes and Wasserman, 2017), anti-competitive patenting strategies (Macdonald, 2004; Blind et al., 2009), the high risk and cost of litigation (Bessen et al., 2011; Bessen and Meurer, 2013; Lemley et al., 2018), and many other problems that have been covered in recent years (Hall and Harhoff, 2012; Eckert and Langinier, 2014; Sampat, 2018). We explain how a more decentralised, transparent, and accessible patent records system offers the flexibility required to address or reduce the impact of many of these issues.

We use the word ‘decentralisation’ very broadly in this work. First, *patent data* and the management of this data is currently quite centralised; patent offices hold and organise all patent information, and we generally trust them with this responsibility. However, this 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 can be largely explained by the centralisation of this system and its interaction with public policy (Jaffe and 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 patent offices that is not being leveraged to address these problems—decentralising some parts of the patenting process may seem a sensible design.

Any discussion concerning reforms that may fix these inefficiencies 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 and 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 may arise.

It has also become clear in recent years that it is not merely inefficiencies within patent *offices* that lie at the root of these problems, but the *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 are embedded. 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 almost entirely out of the control of patent offices (Jaffe and Lerner, 2004; Rai, 2008; Sohn, 2011).^{1,2}

The structure of this document is as follows. We first introduce in Section 2 the fundamental ideas behind DLTs and their use. The main sections of the paper, Sections 3 and 4, split potential 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 concludes, and suggests some avenues for future work.

2 Distributed Ledger Technologies

It is first important to understand, at a high level, how a distributed ledger functions. However, it is also 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. In order to give a general overview,

¹Of 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.

²In the United States, however, there have been recent developments towards increasing patent office autonomy (Wasserman, 2012).

however, what follows is a description of the most common form of distributed ledgers as the technology exists today.

2.1 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, or node, with access to the network is assigned a private key and a public key. 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.³ Every transaction that an entity 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 and 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).⁴

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 in the network 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, most 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 et al., 2017).

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,

³Within the context of this work, a ‘transaction’ generally refers to an interaction between an applicant or patentee and the patent office or examiner.

⁴DLTs, 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.

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 popular blockchains, such as the Bitcoin blockchain or Ethereum, require a considerable amount of computational work to validate blocks of transactions (Krause and 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). Permissioned DLTs, on the other hand, offer much more computationally efficient validation and consensus mechanisms that are entirely different to proof-of-work. This efficiency is possible because the identities of the validators are known and they are generally few (Gramoli, 2017). Therefore, while complexity and computational efficiency are inarguably higher for permissioned DLTs than for traditional systems, we consider these particular downsides to be negligible when considering the implementation of a permissioned DLT.

Blockchain is not the only type of DLT in existence, but it was the first and is the most widely used. However, there exists an entire ecosystem of DLTs that differ from each other in many fundamental ways while sharing a common goal, namely, to provide a decentralised source of information. 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 and Pahl, 2018; Yang et al., 2018; Pervez et al., 2018).⁵ While these specific properties may not be particularly useful for application to patent systems as described in the remainder of this work, we will continue to use the general term DLT so as not to limit our proposals to a single technical framework. After all, the most suitable system may not yet exist.

2.2 On the Necessity of DLTs

The remainder of this work will go into some detail about how DLTs may be used to improve the functioning of current patent systems. However, DLTs are not strictly necessary for any single proposal presented here. So then, why use DLTs at all?

It may be argued that DLTs are only necessary when it is desirable to reduce the reliance on trust between transacting entities or, if possible, do away with trust altogether. This way of thinking could lead one to the conclusion that DLTs are probably unnecessary for our proposals. After all, centralised systems based on mature technological solutions may be able to perform these roles just as well. Inventors and firms already trust the patent office with some of their most valuable assets: their ideas. The question usually associated with these concerns is: Do you need a blockchain? (Wüst and 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?

To answer this question, we must acknowledge that while trust is generally not a problem for developed countries with modern patent systems, trustworthy systems may not always be easily accessible to inventors and firms in less developed regions. However, we will steel-man the argument above by assuming that DLT-based patent record systems are implemented in developed countries, and that inventors deeply trust patent offices with their valuable Intellectual Property (IP). Is there still a reason to use a DLT in this scenario?

Core to the arguments presented in this work is the assertion 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. Some of these possibilities may indeed involve interactions between parties that do not necessarily trust each other such as contracting out patent examination (see Section 4.2). These solutions may, therefore, be thought of as infeasible at present due to the sensitive nature of the information contained in patent applications. Patent-related interactions between distrustful parties may also occur outside the patent office. For example, patents are tradable or licensable assets (Arora and Gambardella, 2010); DLT-based patent information, with

⁵For example, directed acyclic graph-based DLTs (Benčić and Žarko, 2018) are sometimes touted as the natural successors of blockchains (Popov, 2016; Pervez et al., 2018), and are already gaining traction as an appropriate form of DLT for applications such as microtransactions (Popov, 2016).

government-verified ownership and legal status, may reduce uncertainty and other frictions in the markets for these assets. In short, while societal trust in the patent office is not generally a concern, the existence of DLTs means that issues of trust need not restrict our thinking about improvements to patent policy. Additionally, it is also in the interest of governments to remove barriers to innovation and trade that would otherwise benefit society at large—it may be worth adopting DLT-based records solely to encourage more activity in adversarial environments outside the patent office.

It is also noteworthy that firms currently use DLT-based records internally, which suggests that trust is not the only factor driving the adoption of this technology. The benefits of DLTs for streamlining asset management (Savelyev, 2018), data sharing (Xia et al., 2017), and identity management (Dunphy and Petitcolas, 2018), for example, are well-known. In the context of IP, these benefits may be recast as: automation of knowledge asset management through smart contracts or off-ledger code,⁶ a pathway towards a global patent and application database with standardised formats, and transparent and automatic recording of ownership or licensing information. Most of these applications have little to do with distrust between entities, and yet clearly represent a range of possibilities that are either administratively inefficient to operate or otherwise infeasible if implemented within current systems. We believe such benefits present a compelling case for the consideration of DLT-based patent records, in addition to the applications above that are unlocked by a reduced need for trust.

By understanding the advantages of DLTs over traditional means of record-keeping, we can re-frame the primary functions of government-run patent systems to leverage these advantages. What follows is a reimagining of the patent system in a way that offers decentralised, immutable, and transparent records.

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 patents 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

⁶Smart contracts are automatically executing if/else statements that are recorded on a DLT and may or may not be associated with an actual contract—some new DLT infrastructures even distance from this terminology, e.g., in Hyperledger Fabric (Androulaki et al., 2018) smart contracts are called “chaincode.” In the current work, the only distinction we make between smart contracts and (simple) traditional programs is whether or not it is encoded on the DLT (and is therefore immutable and public). Both types of code can, in theory, react to events on or off the DLT.

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 chapter, 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 the patent office (by facilitating policy experimentalism) 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 made 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,⁷ while the European Patent Office (EPO) achieves much the same level of transparency through the European Patent Register.⁸ 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,⁹ 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 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.¹⁰ 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

⁷<https://portal.uspto.gov/pair/PublicPair> (last accessed 07/11/2019)

⁸<https://register.epo.org/regviewer> (last accessed 07/11/2019)

⁹Cryptographic 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.

¹⁰Knowing 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.

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.¹¹

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 and 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).¹² 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.¹³ This metadata

¹¹See e.g., [ipwe.com](https://www.ipwe.com) (last accessed 07/11/2019).

¹²This 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 07/11/2019).

¹³Note, however, that this kind of action should be subject to additional security measures, see Section 6.1.3.

would automatically be updated on the patent record. Many firms *currently* report reassignments to the patent office (Graham et al., 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 document, along with any changes in these entities over time.¹⁴ 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.¹⁵

In summary, DLTs would integrate nicely into any policy that would seek to clarify patent ownership or enforcement status at any particular point in time. 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).¹⁶ 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.

¹⁴See <https://www.govinfo.gov/app/details/FR-2014-01-24/2014-01195> for more information (last accessed 07/11/2019).

¹⁵Note 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 favor of a subsequent transfer that does get reported. (See 35 USC §261).

¹⁶Loss 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 6.1.3.

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 (Lemley and Myhrvold, 2007; Arora and Gambardella, 2010). Unlike markets for other goods, negotiations for IP rights are often closed-door ordeals with significant information asymmetries between parties (Gallini and Wright, 1990; Buenstorf and Geissler, 2012).

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 et al., 2001; Sichelman and O’Connor, 2012; de Rassenfosse, 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 et al., 2015).

More generally, licensing has benefits for society at large. Theoretical work suggests that both ex-ante and ex-post licensing facilitates and incentivises sequential innovation (Green and Scotchmer, 1995), therefore aligning the financial interests of firms with social welfare by accelerating technological progress where invention is cumulative. In a similar vein, some empirical work indicates that the common utilisation of cross-licensing agreements may neutralise the adverse effects that patents could have on follow-on innovation (Sampat and Williams, 2019).

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 et al., 2007).¹⁷ Both the USPTO and WIPO facilitate some form of post-grant patent licensing market,¹⁸ and pre-written contracts could modernise these services: the licensor sets simple contractual terms, places the licenses in some central marketplace,¹⁹ and anyone who would like to license the

¹⁷Patent licensing negotiations are often complicated, strategic, and multilateral ordeals. It is therefore unlikely that smart contracts would simplify, or even be feasible for, many forms of patent licensing.

¹⁸E.g., see <https://www.uspto.gov/learning-and-resources/newsletter/inventors-eye/licensing-options> and <https://www3.wipo.int/wipogreen/en/> (both last accessed 07/11/2019).

¹⁹Markets for patent transfers and licenses in a world with DLT-based patent records would inevitably be set up by third parties in any case.

technology could simply buy the contract.²⁰ 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,²¹ 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 (Shapiro, 2000; Lerner and Tirole, 2015).

All of these licensing and transferal 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.

3.2.3 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 (USPTO, 2014; de Rassenfosse, 2018). 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.²² 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 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 (Shapiro, 2000; Fromer, 2016). To incentivise patentees to use VPMs in this way, they could act as evidence of commercialisation or other usage (Vonneuman, 1956; Trimble, 2016) 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

²⁰In this example, the contract would constitute something akin to a subscription.

²¹As 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.

²²A standard format for VPMs would be ideal; currently, there are no formatting requirements.

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 can 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., Jaffe and Lerner, 2004; Bessen and Meurer, 2008).

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 examine how this 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 and 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 A 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.²³ This way, secrecy benefits the public to some extent

²³The arguments against this proposition generally concern with time-to-market. For

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, 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. 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 is briefly discussed in the current context in Appendix D.

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 follow-on innovation (Furman et al., 2018), in a more timely fashion,²⁴ but may also increase the quality of patent applications. 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 patentabil-

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.

²⁴Early disclosure has been studied extensively in both the theoretical (e.g., Aoki and Spiegel, 2009) and empirical literature (e.g., Graham and Hegde, 2015).

ity of an invention in its current state (assuming a patent is necessary) would have two options: develop the invention further to ensure patentability (this option is explored in a related context in Cotropia, 2009), 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—these developments are discussed in Appendix B.

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 the success of these firms as measures of their capabilities and can be used as a signal to investors (Lemley, 2000; Long, 2002; Hsu and Ziedonis, 2008; Veer and Jell, 2012; Häussler et al., 2012; de Rassenfosse, 2012; Conti et al., 2013; Hall, 2019; Farre-Mensa et al., 2019). Immediate disclosure would speed up these fundraising processes, including pre-grant licensing (Hegde and Luo, 2017) and loans (Saidi and Zaldokas, 2017), 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, countries, and time, (Levin et al., 1987; Harabi, 1995; Brouwer and Kleinknecht, 1999; Cohen et al., 2000; Arundel, 2001; Cohen et al., 2002; Blind et al., 2006; Capponi et al., 2019). 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.²⁵ These applications do not need evidence to demonstrate viability (Freilich, 2018; Freilich and 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

²⁵This practice is legal and done in part to illustrate invention utility and preempt potential infringements

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 (Pakes, 1986; Gambardella et al., 2008) 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 C 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 Weak Patents

Recent years have seen increasing concern about the quality of granted patents (Jaffe and Lerner, 2004; Wagner, 2008; de Rassenfosse et al., 2016). Much of this concern has been focused on the perceived proliferation of so-called ‘weak’ patents (Choi, 2005; Lemley and Shapiro, 2005; Farrell and Shapiro, 2008) 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), variable 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 and Lemley, 1998; Bessen and 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 (Hall and Ziedonis, 2001; Chien, 2010), 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 and 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 and Lerner, 2004) has led to the number of patent applications to the office ‘exploding’ in the past 30 years (Hall, 2004; Strandburg et al., 2006), and the USPTO has generally not managed to keep up with this demand (Ackerman, 2011). While all patenting firms are affected by 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., 2019). While the USPTO announced “specific actions [they] have taken to help reduce the backlog” in 2011,²⁶ and the backlog has indeed decreased, the *percentage* of applications that are granted has risen at the same time (Cotropia et al., 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 and 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* for the remainder of this section, and how DLTs may enable this development.

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.²⁷ This concept was first proposed

²⁶<https://www.uspto.gov/learning-and-resources/newsletter/inventors-eye/uspto-year-review-and-look-forward>

²⁷Note that this proposal is distinct from outsourcing prior art searches as practiced in some jurisdictions (Yamauchi and 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.

and explored rigorously by Thomas (2001). The need for such a system has been echoed more recently (Ghafele et al., 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*,²⁸ and assumed that the benefits to the third-party experts would 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 amongst their peers and future employers. We argue that 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. 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.

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.²⁹ If they deem a particular piece of submitted prior art relevant and uses 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.³⁰ These incentives may be further enhanced when a patent

²⁸The 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).

²⁹This queue-based prior art check is already standard practice, with applicant-supplied prior art reviewed first before an independent search by the examiner.

³⁰One can imagine this process being streamlined for large organisations employing

is of particularly high value, as competitors would like to see the scope of these patents narrowed as much as possible. Competitors often know more about the technology and state of the art than do the examiners. Therefore, in a system where examiners are restricted 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.³¹

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 E. 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

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.

³¹This 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.

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 F).

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 (Thomas, 2001; Sampat, 2005). 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.³² Another foul-play mitigation strategy is to ensure that bounty hunters are subject to a robust identity verification procedure during the bounty hunter registration process.

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,³³ 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 and Duffy, 2008) and can thus perform poorly with little risk (Domberger and 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

³²A 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).

³³Note that this concept has been discussed previously in the legal literature (Abramowicz and Duffy, 2008).

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.³⁴

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 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 and Nagaoka, 2015). In the case of the USPTO, patent examiners often leave to become patent attorneys (Drahos, 2010; Tabakovic and 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 technical 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 use the same submitted prior art information for whole patent families, 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.³⁵ This suggests that, at least in the case of less complex technologies, outsourcing prior art search

³⁴This innovation is also likely to occur outside of the contracting firms by specialised service providers.

³⁵There are indications that these effects are driven in part by a broader scope in the prior art search of the contractors.

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.

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 and 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 Organization.³⁶

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 and 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 et al., 2005; Lichtman and Lemley, 2007; Lemley, 2011). These ideas are discussed in more detail in relation to private examination in Appendix G.

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 USPTO examiner—but also provides checks on the performance and reliability of third party examiners, perhaps in addition to random checks (see also Abramowicz and Duffy, 2008). This system also solves, in part, an

³⁶https://www.wipo.int/meetings/en/doc_details.jsp?doc_id=345737 (last accessed 07/11/2019)

important problem with private examination: the potential for collusion between the applicant and the private examiner. At present, while the incentives, resource constraints, and protocols at some patent offices are undoubtedly skewed in favour of granting patents (Schuett, 2013; Frakes and Wasserman, 2015), 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 does nothing to prevent, for example, leniency towards marginal patents in exchange for return business. Of course, one way to mitigate this 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 USPTO does not agree with the result of the private examination, then the *contractor* must pay the USPTO 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 USPTO 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 and 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, firms 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.³⁷ 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 firms to pay maintenance fees commensurate with the present value of the invention to the firm 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

³⁷There 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 (Cornelli and Schankerman, 1999; Scotchmer, 1999; Gans et al., 2004; Baudry and Dumont, 2006; de Rassenfosse and van Pottelsberghe de la Potterie, 2013).

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 piecewise 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.³⁸ 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 examination quality monitoring.³⁹ 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 in-

³⁸It 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.

³⁹Outsourcing in this manner is not unusual. For example, financial auditing is outsourced to large accounting firms in many countries.

validity 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 and 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.⁴⁰

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 are discriminatory with regard to price: 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 et al., 2004).⁴¹

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

⁴⁰These kinds of checks for examined applications are currently not possible at patent offices due to resource constraints (Allison and Hunter, 2006).

⁴¹There 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 and Jaffe, 2018).

how particular types of information is to be stored will not be discussed here;⁴² 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,⁴³ 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 facilitates 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 judgements 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: amongst 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

⁴²See, e.g., Benet (2014) for an example of a popular DLT-based file storage system.

⁴³There 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.

example, to ensure all transactions of a particular type are recorded in one place, whether it be supply chain management (Kshetri, 2018; Saberi et al., 2019), or workflow management (Fridgen et al., 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 and 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 Section 2) 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).⁴⁴

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.⁴⁵ 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 (granted) patents are concerned.

Most jurisdictions have their own assignee or applicant identification systems. However, 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 from competitors and governments.

⁴⁴A 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)

⁴⁵See, e.g., voltrontrade.com, marcopolo.finance, b3i.tech, hashedhealth.com and dlt.mobi, respectively (all last accessed 07/11/2019).

6 Conclusions and Future Work

Taken individually, most, if not all, of the possibilities for improvement of patent systems offered above 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. DLTs not only make many of these ideas administratively feasible, but are also able to be tested in a step-wise, scalable, and very public manner.

Such a system would 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.

It is clear, however, 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 here. The remainder of this section outlines some of these issues.

6.1 Future Work

6.1.1 Legal Questions

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 the Paris Convention for the Protection of Industrial Property), 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 innovation 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 et al., 2012; Remus and Levy, 2017); however, in the case of patent systems, any solutions may be able to be trialled at a smaller scale through current administrative review procedures (e.g., the Patent Trial and Appeal Board in the United States).

6.1.2 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 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 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.⁴⁶

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 and Simcoe, 2019). This prompts further questions regarding field-specific training courses and certifications, regulations around publicising these certifications, maintaining these certifications at regular intervals, or even 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.⁴⁷ These are small details in the grand scheme; however, they highlight the significant upfront cost of implementing such a significant change to a government-run system.

6.1.3 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 can be utilised. 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. There exist many consensus mechanisms that are very robust to attack (Baliga, 2017; Mingxiao et al., 2017; Cachin and Vukolić, 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 (in contrast to an attack on a public proof-of-work based consensus mechanism which simply requires a lot of computational power (Li et al.,

⁴⁶In 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.

⁴⁷In a recent work, Righi and Simcoe (2019) observe lower grant rates for more specialised patent examiners, indicating more stringent examination.

2017)). 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. Independent of the patent DLT, the use of DLTs for identity management is an active area of research (Shrier et al., 2016; Kuperberg, 2019; Lesavre et al., 2019); these may present much more secure modes of transaction than existing DLTs currently offer.

6.1.4 Theoretical Challenges

Even small modifications to the structure and function of complex social systems, particularly those which 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 open 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.

6.1.5 The Future of IP Policy

While the majority of the concerns above might be considered barriers to the implementation of the current proposals, there are also important questions about the future of IP that will inform the design of this system. For example, how might we build-in interoperability between patent systems and other forms of innovation-incentivising policies such as patent box regimes, prizes or R&D tax credits? 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 all these questions will be an endeavour stretching across disciplinary boundaries.

Acknowledgements

The authors are grateful to Marco Barulli, Emma Francis, Alan Marco, Geoff Sadler, Kees Schuller, Claudia Tapia, Adam Jaffe, Joachim Henkel, Dominique Guellec and Beth Webster for useful discussions.

References

- 112th Congress of the United States of America. Pub. l. no. 112-29, 125 stat. 284, 2011.
<https://www.govinfo.gov/app/details/PLAW-112publ29>.
- 35 U.S.C. § 17 (United States). Invention Secrecy Act of 1951 (2006).
<https://www.govinfo.gov/app/details/USCODE-2011-title35/USCODE-2011-title35-partII-chap17>.
- M. Abramowicz. The danger of underdeveloped patent prospects. *Cornell Law Review*, 92:1065, 2006.
- M. Abramowicz and J. F. Duffy. Ending the patenting monopoly. *University of Pennsylvania Law Review*, 157:1541, 2008.
- L. J. Ackerman. Prioritization: Addressing the patent application backlog at the united states patent and trademark office. *Berkeley Technology Law Journal*, 26(1):67–92, 2011.
- J.-W. Ahn, M. D. Chang, R. Kokku, and P. Watson. Blockchain for open scientific research, June 11 2019. US Patent App. 10/320,574.
- J. R. Allison and S. D. Hunter. On the feasibility of improving patent quality one technology at a time: The case of business methods. *Berkeley Technology Law Journal*, 21:729, 2006.
- J. R. Allison and M. A. Lemley. Empirical evidence on the validity of litigated patents. *AIPPLA Quarterly Journal*, 26:185, 1998.
- J. J. Anderson. Nontechnical disclosure. *Vanderbilt Law Review*, 69:1573, 2016.
- N. P. Anderson. Striking a balance: The pursuit of transparent patent ownership. *Berkeley Technology Law Journal*, 30(4):395–444, 2015.
- E. Androulaki, A. Barger, V. Bortnikov, C. Cachin, K. Christidis, A. De Caro, D. Enyeart, C. Ferris, G. Laventman, Y. Manevich, et al. Hyperledger fabric: a distributed operating system for permissioned blockchains. In *Proceedings of the Thirteenth EuroSys Conference*, page 30. ACM, 2018.
- R. Aoki and Y. Spiegel. Pre-grant patent publication and cumulative innovation. *International Journal of Industrial Organization*, 27(3):333–345, 2009.
- A. Arora and A. Gambardella. Ideas for rent: An overview of markets for technology. *Industrial and Corporate Change*, 19(3):775–803, 2010.
- A. Arora, A. Fosfuri, and A. Gambardella. Markets for technology and their implications for corporate strategy. *Industrial and Corporate Change*, 10(2):419–451, 2001.
- A. Arundel. The relative effectiveness of patents and secrecy for appropriation. *Research Policy*, 30(4):611–624, 2001.
- A. Baliga. Understanding blockchain consensus models. Persistent Systems, 2017.
- S. Bartling. Blockchain for science and knowledge creation. In *Gesundheit digital*, pages 159–180. Springer, 2019.

- M. Baudry and B. Dumont. Patent renewals as options: Improving the mechanism for weeding out lousy patents. *Review of Industrial Organization*, 28(1):41–62, 2006.
- F. M. Benčić and I. P. Žarko. Distributed ledger technology: Blockchain compared to directed acyclic graph. In *2018 IEEE 38th International Conference on Distributed Computing Systems (ICDCS)*, pages 1569–1570. IEEE, 2018.
- J. Benet. IPFS - content addressed, versioned, P2P file system. *arXiv preprint arXiv:1407.3561*, 2014.
- J. Bessen and M. J. Meurer. *Patent failure: How judges, bureaucrats, and lawyers put innovators at risk*. Princeton University Press, 2008.
- J. Bessen and M. J. Meurer. The patent litigation explosion. *Loyola University Chicago Law Journal*, 45:401, 2013.
- J. Bessen, J. Ford, and M. J. Meurer. The private and social costs of patent trolls. *Regulation*, 34:26, 2011.
- K. Blind, J. Edler, R. Frietsch, and U. Schmoch. Motives to patent: Empirical evidence from Germany. *Research Policy*, 35(5):655–672, 2006.
- K. Blind, K. Cremers, and E. Mueller. The influence of strategic patenting on companies’ patent portfolios. *Research Policy*, 38(2):428–436, 2009.
- M. Boldrin and D. K. Levine. *Against intellectual monopoly*. Cambridge University Press Cambridge, 2008.
- M. Boldrin and D. K. Levine. The case against patents. *Journal of Economic Perspectives*, 27(1):3–22, 2013.
- E. Brouwer and A. Kleinknecht. Innovative output, and a firm’s propensity to patent: An exploration of CIS micro data. *Research Policy*, 28(6):615–624, 1999.
- G. Buenstorf and M. Geissler. Not invented here: Technology licensing, knowledge transfer and innovation based on public research. *Journal of Evolutionary Economics*, 22(3), 2012.
- J. E. Cabral, A. Chavan, T. M. Clarke, and J. Greacen. Using technology to enhance access to justice. *Harvard Journal of Law & Technology*, 26:241, 2012.
- C. Cachin and M. Vukolić. Blockchain consensus protocols in the wild. *arXiv preprint arXiv:1707.01873*, 2017.
- G. Capponi, P. Criscuolo, A. Martinelli, and A. Nuvolari. Profiting from innovation: Evidence from a survey of Queen’s Awards winners. *Structural Change and Economic Dynamics*, 49:155–169, 2019.
- C. Catalini and J. S. Gans. Some simple economics of the blockchain. *National Bureau of Economic Research*, (w22952), 2016.
- C. V. Chien. From arms race to marketplace: The complex patent ecosystem and its implications for the patent system. *Hastings Law Journal*, 62:297, 2010.
- J. P. Choi. Live and let live: A tale of weak patents. *Journal of the European Economic Association*, 3(2-3):724–733, 2005.

- W. M. Cohen, R. R. Nelson, and J. P. Walsh. Protecting their intellectual assets: Appropriability conditions and why us manufacturing firms patent (or not). *National Bureau of Economic Research*, (w7552), 2000.
- W. M. Cohen, A. Goto, A. Nagata, R. R. Nelson, and J. P. Walsh. R&D spillovers, patents and the incentives to innovate in Japan and the United States. *Research Policy*, 31(8-9):1349–1367, 2002.
- A. Conti, J. Thursby, and M. Thursby. Patents as signals for startup financing. *The Journal of Industrial Economics*, 61(3):592–622, 2013.
- F. Cornelli and M. Schankerman. Patent renewals and R&D incentives. *RAND Journal of Economics*, pages 197–213, 1999.
- C. A. Cotropia. The folly of early filing in patent law. *Hastings Law Journal*, 61: 65, 2009.
- C. A. Cotropia, C. D. Quillen Jr, and O. H. Webster. Patent applications and the performance of the US Patent and Trademark Office. *Federal Circuit Bar Journal*, 23:179, 2013.
- J. Dai and M. A. Vasarhelyi. Toward blockchain-based accounting and assurance. *Journal of Information Systems*, 31(3):5–21, 2017.
- G. de Rassenfosse. How SMEs exploit their intellectual property assets: Evidence from survey data. *Small Business Economics*, 39(2):437–452, 2012.
- G. de Rassenfosse. Notice failure revisited: Evidence on the use of virtual patent marking. *National Bureau of Economic Research*, (w24288), 2018.
- G. de Rassenfosse and A. B. Jaffe. Are patent fees effective at weeding out low-quality patents? *Journal of Economics & Management Strategy*, 27(1):134–148, 2018.
- G. de Rassenfosse and B. van Pottelsberghe de la Potterie. The role of fees in patent systems: Theory and evidence. *Journal of Economic Surveys*, 27(4):696–716, 2013.
- G. de Rassenfosse, A. B. Jaffe, and E. Webster. Low-quality patents in the eye of the beholder: Evidence from multiple examiners. *National Bureau of Economic Research*, (w22244), 2016.
- S. Domberger and P. Jensen. Contracting out by the public sector: Theory, evidence, prospects. *Oxford Review of Economic Policy*, 13(4):67–78, 1997.
- P. Drahos. *The global governance of knowledge: Patent offices and their clients*. Cambridge University Press, 2010.
- P. Dunphy and F. A. Petitcolas. A first look at identity management schemes on the blockchain. *IEEE Security & Privacy*, 16(4):20–29, 2018.
- A. Eckert and C. Langinier. A survey of the economics of patent systems and procedures. *Journal of Economic Surveys*, 28(5):996–1015, 2014.
- N. El Ioini and C. Pahl. A review of distributed ledger technologies. In *OTM Confederated International Conferences “On the Move to Meaningful Internet Systems”*, pages 277–288. Springer, 2018.

- J. Farre-Mensa, D. Hegde, and A. Ljungqvist. What is a patent worth? Evidence from the U.S. Patent “lottery”. *Journal of Finance*, Forthcoming, 2019.
- J. Farrell and C. Shapiro. How strong are weak patents? *American Economic Review*, 98(4):1347–69, 2008.
- R. Feldman. Plain language patents. *Texas Intellectual Property Law Journal*, 17: 289, 2008.
- M. D. Frakes and M. F. Wasserman. Does the US Patent and Trademark Office grant too many bad patents: Evidence from a quasi-experiment. *Stanford Law Review*, 67:613, 2015.
- M. D. Frakes and M. F. Wasserman. 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, 2017.
- M. D. Frakes and M. F. Wasserman. Irrational ignorance at the patent office. *Vanderbilt Law Review*, 72, 2019.
- J. Freilich. Prophetic patents. 2018.
Available at SSRN: <https://ssrn.com/abstract=3202493>.
- J. Freilich and L. L. Ouellette. Science fiction: Fictitious experiments in patents. *Science*, 364(6445):1036–1037, 2019.
- G. Fridgen, S. Radszuwill, N. Urbach, and L. Utz. Cross-organizational workflow management using blockchain technology-towards applicability, auditability, and automation. In *Proceedings of the 51st Hawaii International Conference on System Sciences*, 2018.
- J. C. Fromer. Patent disclosure. *Iowa Law Review*, 94:539, 2008.
- J. C. Fromer. Dynamic patent disclosure. *Vanderbilt Law Review*, 69:1715, 2016.
- J. L. Furman, M. Nagler, and M. Watzinger. Disclosure and subsequent innovation: Evidence from the patent depository library program. *National Bureau of Economic Research*, (w24660), 2018.
- N. T. Gallini and B. D. Wright. Technology transfer under asymmetric information. *The RAND Journal of Economics*, pages 147–160, 1990.
- A. Gambardella, P. Giuri, and A. Luzzi. The market for patents in Europe. *Research Policy*, 36(8):1163–1183, 2007.
- A. Gambardella, D. Harhoff, and B. Verspagen. The value of European patents. *European Management Review*, 5(2):69–84, 2008.
- J. S. Gans, S. P. King, and R. Lampe. Patent renewal fees and self-funding patent offices. *Topics in Theoretical Economics*, 4(1), 2004.
- J. S. Gans, D. H. Hsu, and S. Stern. The impact of uncertain intellectual property rights on the market for ideas: Evidence from patent grant delays. *Management Science*, 54(5):982–997, 2008.
- R. Ghafele, B. Gibert, and P. DiGiammarino. Driving innovation through patent application review: The power of crowdsourcing prior art search. *Journal of Intellectual Property Rights*, pages 303–308, 2011.

- S. Graham and D. Hegde. Disclosing patents' secrets. *Science*, 347(6219):236–237, 2015.
- S. J. Graham, A. C. Marco, and A. F. Myers. Patent transactions in the marketplace: Lessons from the USPTO patent assignment dataset. *Journal of Economics & Management Strategy*, 27(3):343–371, 2018.
- V. Gramoli. From blockchain consensus back to Byzantine consensus. *Future Generation Computer Systems*, 2017.
- J. R. Green and S. Scotchmer. On the division of profit in sequential innovation. *The RAND Journal of Economics*, pages 20–33, 1995.
- B. H. Hall. Business method patents, innovation, and policy. *National Bureau of Economic Research*, (w9717), 2003.
- B. H. Hall. Exploring the patent explosion. *Journal of Technology Transfer*, 30(1-2):35–48, 2004.
- B. H. Hall. Is there a role for patents in the financing of new innovative firms? *Industrial and Corporate Change*, 28(3):657–680, 2019.
- B. H. Hall and D. Harhoff. Recent research on the economics of patents. *Annual Review of Economics*, 4(1):541–565, 2012.
- B. H. Hall and R. H. Ziedonis. The patent paradox revisited: An empirical study of patenting in the US semiconductor industry, 1979-1995. *RAND Journal of Economics*, pages 101–128, 2001.
- B. H. Hall, C. Helmers, and G. Von Graevenitz. Technology entry in the presence of patent thickets. *National Bureau of Economic Research*, (w21455), 2015.
- N. Harabi. Appropriability of technical innovations an empirical analysis. *Research Policy*, 24(6):981–992, 1995.
- C. Häussler, D. Harhoff, and E. Müller. To be financed or not...: The role of patents for venture capital-financing. *ZEW-Centre for European Economic Research Discussion Paper*, (09-003), 2012.
- D. Hegde and H. Luo. Patent publication and the market for ideas. *Management Science*, 64(2):652–672, 2017.
- D. J. Hemel and L. L. Ouellette. Innovation policy pluralism. *Yale Law Journal*, 128:544, 2019.
- J. Henkel and H. Zischka. How many patents are truly valid? extent, causes, and remedies for latent patent invalidity. *European Journal of Law and Economics*, 48(2):195–239, 2019.
- R. T. Holden and A. Malani. Can blockchain solve the hold-up problem in contracts? *National Bureau of Economic Research*, (w25833), 2019.
- D. H. Hsu and R. H. Ziedonis. Patents as quality signals for entrepreneurial ventures. In *Academy of Management Proceedings*, number 1 in 2008, pages 1–6. Academy of Management, 2008.
- L. Hurwicz. The design of mechanisms for resource allocation. *American Economic Review*, 63(2):1–30, 1973.

- IP Australia. Patent manual of practice & procedure, 2018.
http://manuals.ipaustralia.gov.au/patents/national/examination/2.13.11.1_Notifications.htm, Retrieved 30/08/2019.
- A. B. Jaffe and J. Lerner. *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do about it*. Princeton University Press, 2004.
- J. A. Jeffery. Preserving the presumption of patent validity: An alternative to outsourcing the US patent examiner’s prior art search. *Catholic University Law Review*, 52:761, 2002.
- P. H. Jensen and R. E. Stonecash. Incentives and the efficiency of public sector-outsourcing contracts. *Journal of Economic Surveys*, 19(5):767–787, 2005.
- R. Jensen and M. Thursby. Proofs and prototypes for sale: The licensing of university inventions. *American Economic Review*, 91(1):240–259, 2001.
- X. Jin, S. Spangler, Y. Chen, K. Cai, R. Ma, L. Zhang, X. Wu, and J. Han. Patent maintenance recommendation with patent information network model. In *2011 IEEE 11th International Conference on Data Mining*, pages 280–289. IEEE, 2011.
- E. R. Kazenske. The future of prior art searching at the United States Patent and Trademark Office. *World Patent Information*, 25(4):283–287, 2003.
- E. W. Kitch. The nature and function of the patent system. *The Journal of Law and Economics*, 20(2):265–290, 1977.
- M. J. Krause and T. Tolaymat. Quantification of energy and carbon costs for mining cryptocurrencies. *Nature Sustainability*, 1(11):711, 2018.
- N. Kshetri. Blockchain’s roles in meeting key supply chain management objectives. *International Journal of Information Management*, 39:80–89, 2018.
- M. Kuperberg. Blockchain-based identity management: A survey from the enterprise and ecosystem perspective. *IEEE Transactions on Engineering Management*, 2019.
- J.-J. Laffont and J. Tirole. *A theory of incentives in procurement and regulation*. MIT press, 1993.
- M. A. Lemley. Reconceiving patents in the age of venture capital. *Journal of Small & Emerging Business Law*, 4:137, 2000.
- M. A. Lemley. Can the patent office be fixed. *Marquette Intellectual Property Law Review*, 15:295, 2011.
- M. A. Lemley. Ready for patenting. *Boston University Law Review*, 96:1171, 2016.
- M. A. Lemley and N. Myhrvold. How to make a patent market. *Hofstra Law Review*, 36:257, 2007.
- M. A. Lemley and B. Sampat. Is the patent office a rubber stamp. *Emory Law Journal*, 58:181, 2008.
- M. A. Lemley and C. Shapiro. Probabilistic patents. *Journal of Economic Perspectives*, 19(2):75–98, 2005.

- M. A. Lemley, D. Lichtman, and B. N. Sampat. What to do about bad patents. *Regulation*, 28(4):10–13, 2005.
- M. A. Lemley, K. Richardson, and E. Oliver. The patent enforcement iceberg. *Texas Law Review*, 97:801, 2018.
- J. Lerner and J. Tirole. Standard-essential patents. *Journal of Political Economy*, 123(3):547–586, 2015.
- L. Lesavre, P. Varin, P. Mell, M. Davidson, and J. Shook. A taxonomic approach to understanding emerging blockchain identity management systems. *arXiv preprint arXiv:1908.00929*, 2019.
- R. C. Levin, A. K. Klevorick, R. R. Nelson, S. G. Winter, R. Gilbert, and Z. Griliches. Appropriating the returns from industrial research and development. *Brookings Papers on Economic Activity*, 1987(3):783–831, 1987.
- X. Li, P. Jiang, T. Chen, X. Luo, and Q. Wen. A survey on the security of blockchain systems. *Future Generation Computer Systems*, 2017.
- D. Lichtman and M. A. Lemley. Rethinking patent law’s presumption of validity. *Stanford Law Review*, 60:45, 2007.
- C. Long. Patent signals. *The University of Chicago Law Review*, pages 625–679, 2002.
- P. O. Long. Invention, authorship, “intellectual property,” and the origin of patents: Notes toward a conceptual history. *Technology and Culture*, 32(4):846–884, 1991.
- S. Macdonald. When means become ends: Considering the impact of patent strategy on innovation. *Information Economics and Policy*, 16(1):135–158, 2004.
- S. J. Marsnik. Will the America Invents Act post-grant review improve the quality of patents? A comparison with the European Patent Office opposition. *The Changing Face of US Patent Law and Its Impact on Business Strategy*, page 183, 2013.
- S. M. McJohn. Patents: Hiding from history. *Santa Clara Computer & High Technology Law Journal*, 24:961, 2007.
- P. S. Menell and M. J. Meurer. Notice failure and notice externalities. *Journal of Legal Analysis*, 5(1):1–59, 2013.
- D. Mingxiao, M. Xiaofeng, Z. Zhe, W. Xiangwei, and C. Qijun. A review on consensus algorithm of blockchain. In *2017 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, pages 2567–2572. IEEE, 2017.
- P. Mirowski. *Science-mart*. Harvard University Press, 2011.
- S. Nakamoto et al. Bitcoin: A peer-to-peer electronic cash system. 2008.
- B. S. Noveck. Peer to patent: Collective intelligence, open review, and patent reform. *Harvard Journal of Law & Technology*, 20:123, 2006.
- L. L. Ouellette. Do patents disclose useful information? *Harvard Journal of Law & Technology*, 25(2):545, 2012.
- L. L. Ouellette. Pierson, peer review, and patent law. *Vanderbilt Law Review*, 69: 1825, 2016.

- J. Owen-Smith and W. W. Powell. The expanding role of university patenting in the life sciences: Assessing the importance of experience and connectivity. *Research Policy*, 32(9):1695–1711, 2003.
- A. Pakes. Patents as options: Some estimates of the value of holding European patent stocks. *Econometrica: Journal of the Econometric Society*, pages 755–784, 1986.
- Patents Act 1977 (United Kingdom). Patents Act 1977, c. 22.
Available at: <https://tinyurl.com/UKPatentSecrecy>.
- H. Pervez, M. Muneeb, M. U. Irfan, and I. U. Haq. A comparative analysis of DAG-based blockchain architectures. In *2018 12th International Conference on Open Source Systems and Technologies (ICOSST)*, pages 27–34. IEEE, 2018.
- S. Popov. The tangle. *IOTA Whitepaper*, 2016.
- F. D. Prager. A history of intellectual property from 1545 to 1787. *Journal of the Patent Office Society*, 26:711, 1944.
- A. K. Rai. Growing pains in the administrative state: The patent office’s troubled quest for managerial control. *University of Pennsylvania Law Review*, 157:2051, 2008.
- D. Remus and F. Levy. Can robots be lawyers: Computers, lawyers, and the practice of law. *Georgetown Journal of Legal Ethics*, 30:501, 2017.
- C. Righi and T. Simcoe. Patent examiner specialization. *Research Policy*, 48(1):137–148, 2019.
- S. Saberi, M. Kouhizadeh, J. Sarkis, and L. Shen. Blockchain technology and its relationships to sustainable supply chain management. *International Journal of Production Research*, 57(7):2117–2135, 2019.
- F. Saidi and A. Zaldokas. How does firms’ innovation disclosure affect their banking relationships?
Available at SSRN: <https://ssrn.com/abstract=2715925>, 2017.
- B. Sampat. Determinants of patent quality: An empirical analysis. 2005.
Available at <https://www.immagic.com/eLibrary/ARCHIVES/GENERAL/COLUMBIA/C050902S.pdf>.
- B. Sampat and H. L. Williams. How do patents affect follow-on innovation? Evidence from the human genome. *American Economic Review*, 109(1):203–36, 2019.
- B. N. Sampat. A survey of empirical evidence on patents and innovation. *National Bureau of Economic Research*, (w25383), 2018.
- A. Savelyev. Copyright in the blockchain era: Promises and challenges. *Computer Law & Security Review*, 34(3):550–561, 2018.
- F. Schuett. Patent quality and incentives at the patent office. *The RAND Journal of Economics*, 44(2):313–336, 2013.
- S. Scotchmer. On the optimality of the patent renewal system. *The RAND Journal of Economics*, pages 181–196, 1999.

- S. B. Seymore. Heightened enablement in the unpredictable arts. *UCLA Law Review*, 56:127, 2008.
- S. B. Seymore. The presumption of patentability. *Minnesota Law Review*, 97:990, 2012.
- C. Shapiro. Navigating the patent thicket: Cross licenses, patent pools, and standard setting. *Innovation Policy and the Economy*, 1:119–150, 2000.
- D. Shrier, W. Wu, and A. Pentland. Blockchain & infrastructure (identity, data security). *Massachusetts Institute of Technology-Connection Science*, 1(3):1–19, 2016.
- T. Sichelman and S. O’Connor. Patents as promoters of competition: The guild origins of patent law in the Venetian Republic. *San Diego Law Review*, 49:1267, 2012.
- J. L. Sohn. Can’t the PTO get a little respect. *Berkeley Technology Law Journal*, 26:1603, 2011.
- D. F. Spulber. Should business method inventions be patentable? *Journal of Legal Analysis*, 3(1):265–340, 2011.
- J. E. Stiglitz. Economic foundations of intellectual property rights. *Duke Law Journal*, 57:1693, 2007.
- K. J. Strandburg, G. Csardi, J. Tobochnik, and P. Erdi. Law and the science of networks: An overview and an application to the patent explosion. *Berkeley Technology Law Journal*, 21:1293, 2006.
- H. Tabakovic and T. G. Wollmann. From revolving doors to regulatory capture? Evidence from patent examiners. *National Bureau of Economic Research*, (w24638), 2018.
- J. R. Thomas. 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*, pages 95–132. Emerald Group Publishing Limited, 2001.
- J. R. Thomas. The responsibility of the rulemaker: Comparative approaches to patent administration reform. *Berkeley Technology Law Journal*, 17:727, 2002.
- L. C. Thurow. Needed: A new system of intellectual property rights. *Harvard Business Review*, 75:94–107, 1997.
- M. Trimble. Patent working requirements: Historical and comparative perspectives. *UC Irvine Law Review*, 6:483, 2016.
- United States Patent and Trademark Office. USPTO 21st century strategic plan. <https://www.uspto.gov/about-us/performance-and-planning/strategic-planning/strategic-plan-21st-century>, 2003.
- United States Senate. 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>, 1992.

- USPTO. Report on virtual marking, 2014.
https://www.uspto.gov/sites/default/files/aia_implementation/VMreport.pdf.
- T. Veer and F. Jell. Contributing to markets for technology? A comparison of patent filing motives of individual inventors, small companies and universities. *Technovation*, 32(9-10):513–522, 2012.
- N. A. Vonneuman. Conditionally exclusive patent rights and the patent clause of the constitution. *American Journal of Comparative Law*, 5:391, 1956.
- R. P. Wagner. Understanding patent-quality mechanisms. *University of Pennsylvania Law Review*, 157:2135, 2008.
- M. F. Wasserman. The changing guard of patent law: Chevron deference for the PTO. *William & Mary Law Review*, 54:1959, 2012.
- WIPO. Administrative instructions under the Patent Cooperation Treaty, 2012.
https://www.wipo.int/pct/en/texts/ai/s801.html#_801_a, Retrieved 30/08/2019.
- K. Wüst and A. Gervais. Do you need a blockchain? In *2018 Crypto Valley Conference on Blockchain Technology (CVCBT)*, pages 45–54. IEEE, 2018.
- Q. Xia, E. B. Sifah, K. O. Asamoah, J. Gao, X. Du, and M. Guizani. Medshare: Trust-less medical data sharing among cloud service providers via blockchain. *IEEE Access*, 5:14757–14767, 2017.
- X. Xu, I. Weber, M. Staples, L. Zhu, J. Bosch, L. Bass, C. Pautasso, and P. Rimba. A taxonomy of blockchain-based systems for architecture design. In *2017 IEEE International Conference on Software Architecture*, pages 243–252. IEEE, 2017.
- I. Yamauchi and S. Nagaoka. Does the outsourcing of prior art search increase the efficiency of patent examination? Evidence from Japan. *Research Policy*, 44(8):1601–1614, 2015.
- W. Yang, S. Garg, A. Raza, D. Herbert, and B. Kang. Blockchain: Trends and future. In *Pacific Rim Knowledge Acquisition Workshop*, pages 201–210. Springer, 2018.
- S. Yelderman. Prior art in inter partes review. *Iowa Law Review*, 104:2705–2734, 2019.

A 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 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.⁴⁸ 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

⁴⁸In 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.

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.

B 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 numerous⁴⁹ 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,⁵⁰ and the state legislature of Vermont in the United States has passed a law allowing blockchain records as evidence in state courts,⁵¹ 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.⁵²

⁴⁹Examples include Bernstein (bernstein.io) and iDefendo (idefendo.com) (all last accessed 07/11/2019).

⁵⁰<https://tinyurl.com/blockchainChina> (last accessed 07/11/2019)

⁵¹<https://legislature.vermont.gov/statutes/section/12/081/01913> (last accessed 07/11/2019)

⁵²Taken 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.

C 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.⁵³ It would additionally disadvantage firms without the resources to fund the development of a functioning prototype before the investment that a patent may attract (Farre-Mensa et al., 2019), or those that intend to license the technology to firms with development capabilities as might universities (Jensen and Thursby, 2001; Owen-Smith and Powell, 2003).

To combat intentionally obscured patent specifications,⁵⁴ 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.⁵⁵ 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

⁵³Enablement 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 (Seymore, 2008; Ouellette, 2012, 2016), which is out of scope of the current work but may also be facilitated by DLT-based records and feedback (Bartling, 2019; Ahn et al., 2019).

⁵⁴This behaviour likely occurs within the current system as well (McJohn, 2007; Feldman, 2008).

⁵⁵If 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).

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).

D 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 (35 U.S.C. § 17 (United States); Patents Act 1977 (United Kingdom)). The checks for inventions of this nature necessarily take place 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’ 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).

E 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.

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

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.⁵⁷ 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.

F 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.⁵⁸

⁵⁷At 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).

⁵⁸See e.g., loci.io, operem.com, blockchainyourip.com, bernstein.io, artifacts.ai, and idefendo.com (all last accessed 07/11/2019).

G Gold-Plated Patents, SPER, and Private Examiner Incentives

One idea to reduce levels of unnecessary patent litigation, popularised by Mark Lemley and colleagues (Lemley et al., 2005; Lichtman and Lemley, 2007; Lemley, 2011) and directed at the USPTO, proposes a weakening of the default presumption of validity and encourages a move towards an *earned* presumption of validity.⁵⁹ This is sometimes framed as the ‘gold-plating’ of a patent (Lemley et al., 2005; Lichtman and 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 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 and 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 classification⁶⁰ that would result in a SPER (Allison and 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

⁵⁹U.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).

⁶⁰Applications with the SPER-enforced class as a *secondary* classification were exempt.

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 review⁶¹—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 (Laffont and Tirole, 1993; Jensen and Stonecash, 2005). 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.

⁶¹ *Both incorrect rejections and incorrect grants*