CrossMark

The Search of Prior Art and the Revelation of Information by Patent Applicants

Corinne Langinier[1] · Philippe Marcoul[2]

Abstract To learn the patentability of an innovation, both applicant and examiner search through the set of related inventions. The applicant searches first and chooses to reveal his findings to the examiner, who performs a complementary search and decides whether to grant a patent. We analyze this process with a model of bilateral search for information. We show that the applicant may strategically conceal information, and the examiner makes her search contingent upon the revealed information. To remedy information concealment, we focus on two mechanisms: a double-review policy and a commitment mechanism. Both mechanisms induce more revelation of information.

Keywords Incentives · Information · Patents · Prior art

1 Introduction

The issuance of patents of questionable validity is a major concern for patent policy. In an effort to improve the quality of the patent prosecution process, the U.S. Patent and Trademark Office (PTO) has launched several initiatives.[1] To contribute to this

---

[1] The patent prosecution describes the interaction between a patent applicant (or his representatives) and the PTO during the patent application process. Examples of initiatives for reforming the U.S. patenting process can be found in The Patent Reform Act of 2011 or the USPTO website.

✉ Corinne Langinier
   langinie@ualberta.ca

   Philippe Marcoul
   marcoul@ualberta.ca

[1] Department of Economics, University of Alberta, 8-14 Tory Building, Edmonton, AB T6G 2H4, Canada

[2] Department of Resource Economics and Environmental Sociology, University of Alberta, 503 General Services Building, Edmonton, AB T6G 2H1, Canada

discussion, we present a theoretical model of the application and granting of a patent. A key aspect of our modeling approach focuses on the actions of the applicant and the examiner as they search through the set of related inventions, which is referred to as *prior art* in the patent jargon.[2]

This search for prior art is crucial as it permits an assessment of the patentability of an innovation. The examiner's own search of prior art may reveal previous similar innovations or an existing piece of knowledge that shows that the innovation under scrutiny is not novel and, hence, should not be granted a patent. Search is important also because it defines the key tension in our model. By letting the applicant withhold information on quality-relevant prior art, we allow the applicant to influence the examiner's search effort. This strategic concealment of prior art impedes the examiner's inference process, and potentially leads to the granting of a patent to a non-patentable innovation.[3]

Formally, the patent granting process is modeled as a sequential bilateral information gathering game. At the outset, the patentability of the innovation is unknown; consequently, both applicant and examiner must search to identify it. The ease with which the examiner conducts her search depends on the prior art that is disclosed by the applicant. Upon discovering abundant prior art, the applicant may indeed conceal some of the prior art that he discovered after he learned that his innovation is not patentable—knowing that strategic withholding increases examiner's screening costs. Then, relying on her own search but also on what the applicant has chosen to reveal, the examiner processes the application and eventually determines the patentability of the innovation.

We show that the examiner's screening intensity is contingent on the amount of prior art that the applicant reports, and that applicants with non-patentable innovations tend to conceal some prior art. Applicants do search for prior art: essentially, because they fear that forgotten prior art might be a cause for future costly prosecution or even invalidation.

Since concealing information is associated with non-patentable innovations, we then focus on mechanisms that enhance the flow of information that is provided by applicants. The first mechanism is a second-review process that is contingent on the amount of prior art that is disclosed. We show that a second review influences the applicant's decision to reveal prior art, but its overall implication is unclear. A reduction of strategic concealment of information in one well scrutinized field might increase the fraction of "bad" innovators who apply for patents in other less-scrutinized fields.

In response to these ambiguous results, we introduce a second mechanism that requires the examiner to commit *ex ante* to pre-defined levels of scrutiny. We show that the optimal scrutiny of the examiner is uniform across all applications, which induces full revelation of prior art.

---

[2] From the early decision to file for a patent to its final granting, most of the patenting process involves finding, reporting, and checking prior art. Ultimately, the quality of the application combined with the examiner's search efforts determines the outcome of the patenting process.

[3] In practice, examiners spend an average of only 8–25 h on each application (Federal Trade Commission 2003), which suggests that patent screening may be inadequate in identifying the patentability of the innovation.

The paper is organized as follows: in Sect. 2, we present the related literature and the details of the patent filing and prosecution process. The model is presented in Sect. 3. In Sect. 4, we describe the technologies of examination and prior art search. In Sect. 5, we analyze the non-commitment case. We analyze the second-review policy in Sect. 6, and the commitment policy in Sect. 7. We conclude and derive policy implications in Sect. 8.

2 Related Literature, Patent Filing, and the Prosecution Process

Although there is an abundant patent literature,[4] only a few contributions study the process by which patents are granted (Farrell and Merges 2004; Kesan 2002; Merges 1999).[5]

The patent-granting process is standard: When filing a patent, an applicant provides information that contains references to previous innovations and patents upon which the innovation improves or from which it diverges. The nature of this information is usually diverse, and any relevant piece of evidence (theses or scientific articles) constitutes prior art. In the U.S. applicants are not legally required to search for prior art; but if they do, they must disclose it (a "duty of candor").

In addition, according to the "doctrine of inequitable conduct", applicants should not disclose false information. However, the PTO has made it clear that applications will not be investigated and rejected based solely on a violation of the duty to disclose prior art (Kesan 2002).

To assess the patentability of an innovation, an examiner relies on the content of the application and her own prior art search.[6] The returns to search effort for the applicant and for the examiner are a function of the amount of prior art, which can vary across technological fields.

In more mature fields (e.g., pharmaceuticals) prior art is abundant and easy to find. In emerging fields, prior art is less accessible as it is more likely to be disseminated across scientific publications or other sources.[7] Overall, the information that is revealed by applicants should be a function of their search efforts and of the amount of existing prior art.

Innovators' incentives to search for prior art before and after undertaking any R&D investments have been studied by Atal and Bar (2010). In their model, an applicant who did not find invalidating prior art at the initial research stage will be

---

[4] The patent literature focuses on litigation and settlement (Lanjouw and Schankerman 2001; Crampes and Langinier 2002), or on the impact of rules on patent value (Chang 1995; Scotchmer 1996; O'Donoghue 1998).

[5] A few contributions investigate the patent examination process, but they do not consider the applicant's search for prior art (Chiou 2008; Langinier and Marcoul 2014). By modeling the patenting process, Picard et al. (2013) consider that the quality of granted patents depends on the PTO's examination effort. However, they do not model the applicant's incentives to distort the examination process.

[6] An innovation is patentable if it is novel, non-obvious, and useful.

[7] In these fields, prior art is usually not in the form of a patent (Allison and Hunter 2006, and Sampat 2010). For instance, an unpublished thesis can represent prior art that may be used to invalidate a patent. However, finding this information might prove difficult for the examiner. On a similar note, Jaffe and Lerner write "examiners are not very good at finding non-patented prior art" (Jaffe and Lerner 2004).

optimistic regarding the quality of his innovation at the patenting stage. Thus, he will be reluctant to search further. However, they do not analyze the strategic revelation of information or the examiner's behavior.

Although an applicant's not searching for prior art (and concealing it) will hinder the examiner's ability to reject the application, it is also costly as patents with missing prior art tend to be invalidated in court (Allison and Lemley 1998). Hence, when deciding to search, an applicant must strike a balance between strengthening his (potentially good) patents and obtaining (ultimately bad) patents. In some fields, patents tend to have more prior art cited, which is an indication that search has taken place.

Due to the lack of accessible data, only recently has empirical attention been focused on prior art problems and on opening the "black box" of the patent-granting process (Cockburn et al. 2003). The findings of Sampat (2010), and Alcácer et al. (2009) suggest that applicants and examiners do not necessarily have the same information, or the same incentives to search and reveal information. Examiners seem to be less informed than are applicants about prior art, and may face particular challenges in searching for it (Jaffe and Lerner 2004).

Prior art that is provided by applicants can be sparse,[8] even though the relevant information exists and can be accessed (albeit at substantial cost).[9] Lampe (2012) finds empirical evidence that applicants conceal prior art, which validates our main finding that they have an incentive not to reveal all of it.

Patent reforms have been suggested to improve the U.S. patent system, such as an opposition system[10] (Merges 1999) or a two-tiered system in which applicants can choose between a regular patent or a more expensive 'gold plated' patent with a more rigorous examination and an enhanced presumption of validity (Lemley et al. 2005; Atal and Bar 2014). Even if it is important to provide better incentives for applicants to find and disclose prior art (Farrell and Merges 2004), so far none of these proposals are concerned with such incentives.

Our paper also relates to the auditing and monitoring literature: Contracts must sometimes reward agents for announcing bad news (Levitt and Snyder 1997); the principal commits *ex ante* to an inefficient *ex post* outcome. In our setting, the examiner may structure her commitment to induce applicants to search and fully reveal prior art. It is similar to a delegation model where the principal commits to not interfere in the agent's activities to encourage initiative (Aghion and Tirole 1997). Without a credible mechanism for committing to an audit, there is no equilibrium in which only good applicants apply for patents.[11]

---

[8] Alcácer and Gittelman (2006) show that, during 2001–2003, 40 % of patents had all citations listed by examiners. Examiners tend to cite patents that are owned by the applicants themselves (Sampat 2010). As it is unlikely that applicants are unable to find their own patents, this suggests that many applicants do not search for, or fail to disclose, prior art. Some financial patents seem to have clearly incomplete prior art (Lerner 2002).

[9] There exist websites where prior art can be added to a patent (e.g., PatentFizz). For instance, on Bountyquest.com, one can post a request for prior art for a reward from $10,000 to $30,000.

[10] For details on the European opposition system, see Graham et al. (2002) and Friebel et al. (2006).

[11] This is related to the literature on optimal contracts with no commitment to an audit policy (Khalil 1997; Strausz 2005; Mitusch 2006).

3 The Model

We consider a sequential game in which a patent applicant is endowed with an innovation and an examiner must handle a patent application. The applicant searches for prior art and, depending on his effort and the scarcity of prior art, he discovers one of two exogenous levels of prior art: either $x_h \equiv 1$, or $x_l < 1$. Given what he has discovered, he then decides the amount of prior art $\tilde{x}$ to reveal in the application (either $\tilde{x} = 1$ or $\tilde{x} = x_l$). When she receives the application, the examiner processes it, checks the application's references and performs her own prior art search that may (or may not) reveal similar existing innovations. Based on the information that she discovers, she then decides whether to grant a patent. Below, we describe each of these stages in detail.

At the outset, neither the applicant nor the examiner knows the patentability and how much prior art exists. With probability p, it is a patentable (or "good") innovation, and with probability $(1 - p)$ it is a non-patentable (or "bad") innovation.[12] The unconditional probability that the prior art is abundant (resp., scarce) is γ (resp., $1 - \gamma$). *Ex ante*, γ and p are common knowledge to the applicant and the examiner.[13]

Before filing for a patent, the applicant chooses his prior art search effort. By exerting an effort $e \in [0, 1]$, which has a disutility $c(e)$, he finds $x \in X = \{x_l, 1\}$. We assume that $c'(e) > 0$, $c''(e) > 0$, $c(0) = c'(0) = 0$, $\lim_{e \to 1} c(e) = \lim_{e \to 1} c'(e) = \infty$, and $c'''(e) > 0$.[14]

His effort generates a probability of finding prior art conditional upon it being abundant or scarce. We assume that this conditional distribution is uniform. The prior art found by the applicant can take two values: x_l or x_h, where $0 < x_l < x_h \equiv 1$, which is known *ex ante* by both the examiner and the applicant. To simplify, we rule out the possibility of revealing no prior art. We implicitly assume that the examiner does not grant patents to applications that have empty prior art content.

When the prior art is abundant, the applicant finds 1 with probability $prob(x = 1 \,|\text{prior art is abundant}) = e$. When it is scarce, he never finds 1 as $prob(x = 1 \,|\text{prior art is scarce}) = 0$ but he finds x_l with probability $prob(x = x_l \,|\text{prior art is scarce}) = 1$. If $e > 0$, the applicant learns whether his innovation is patentable (or not) when he finds the whole existing prior (i.e., when $x = 1$ and the prior art is abundant, and when $x = x_l$ and the prior art is scarce).[15,16] The discovery of the patentability is soft information, and its content might be

---

[12] All parameters and variables of the game are summarized in Table 1.

[13] By abundant, we mean that, in certain technological areas, there are numerous patented prior art references that are easily accessible in patent data sets. Conversely, in other technological areas, prior art is scarce and searches in existing patent data sets will reveal very few references to existing innovations.

[14] This last assumption is not necessary for our results, but it simplifies the calculations.

[15] We assume perfect learning by innovators but we could allow this learning to be interpreted as a "hint" that some innovations are better than others. We would obtain similar qualitative results by assuming that a fraction of applicants learn their type. This aspect underlies the fact that a fair amount of learning about novelty takes place when applicants prepare their applications (Trajtenberg et al. 2000).

[16] Whether the applicant learns that the prior art comes from an abundant field or not is irrelevant.

omitted without altering the amount submitted, even though this is against the law.[17] If $e = 0$, he does not learn the patentability of his innovation.

After having searched for prior art, the applicant applies for a patent and reveals $\widetilde{x}$.[18] He can conceal some prior art (reveal x_l instead of 1), but can never forge an application that would contain more prior art than he actually found. This aspect captures the possibility that applicants can withhold citations to relevant prior art.

From the examiner's perspective, an applicant's revelation of x_l can arise from four possible circumstances: *a*) The applicant has chosen not to search (i.e., $e = 0$); *b*) The innovation is in a sparse prior art field; *c*) The innovation is in an abundant area, but the applicant only found x_l; or *d*) The applicant found 1, but he also learned that his innovation is not patentable and strategically decided to withhold some of the prior art.

The examiner does not observe the applicant's effort, nor does she observe what he actually found. She has prior beliefs about the patentability of the innovation and the field to which it belongs, and she must exert an effort to establish its patentability. In practice, this effort involves a mix of both checking the prior art that the applicant has reported and performing her own search of prior art. Arguably, searching for prior art tends to be costlier than merely checking prior art references.

We therefore assume that the marginal cost of this (composite) effort is higher when the patenting process involves more 'pure' search and fewer prior art references are transmitted to the examiner. More precisely, the examiner receives an application with prior art $\widetilde{x}$, updates her beliefs about patentability accordingly,[19] and exerts a search effort $E \in [0, 1]$ at cost

$$C(E, \widetilde{x}) = \frac{1}{2\widetilde{x}} E^2. \tag{1}$$

Her effort generates a probability E to discover whether the innovation is patentable. With probability $(1 - E)$, she does not discover invalidating information.[20] Specification (1) implies that the examiner's search cost is higher when the prior art reported $\widetilde{x}$ is small. Our goal is to model the "complementarity" between

---

[17] If an applicant finds a citation that might invalidate his patent, he may remove it and pretend he never encountered it, even though this is illegal. Alternatively, we could think about the quality of disclosure: The applicant might disclose prior art in a way that is not very helpful to the examiner.

[18] For simplicity, we assume that the examiner always applies for a patent, since it is costless. Patenting decisions have already been studied by Atal and Bar (2010).

[19] She does not update her beliefs about the field, since what matters is to grant a patent to a patentable innovation.

[20] Admittedly, our search process is simple to allow for an overall tractable model. In a less parsimonious model, the examiner's search and learning process could happen in several stages where current discoveries about prior art influence future search effort (and, hence, future discoveries), which would create many possible search pathways and give rise to stopping rules. In this setting, the examiner would stop searching when she discovers sufficiently similar innovations since it would essentially show that the innovation does not meet the patentability requirements. For other search pathways and for convex sampling costs, the examiner would abandon search when the sampling of prior art sources keeps giving her anything but prior art close to the innovation described in the application.

the information that is provided by the applicant and the examiner's efficiency.[21] The functional form (1) captures this aspect of the patenting process.[22]

Based on the information that is provided by the applicant and the prior art she has discovered, the examiner grants a patent or not. Establishing the patentability depends on the quality of the information cues (e.g., innovation attributes, prior art) that have been provided by the applicant, the sources that are used by the examiner (e.g., the database), and her ability to perform the examination.[23] When receiving an application containing 1, the examiner will check the prior art cited and evaluate claims against the existing prior art that 'surrounds' the innovation. Thus, in this case, the prior art search is likely to be minimal. If "x_l" has been reported to her, she will have to search for additional prior art (that may have been concealed); and, if additional prior art is found, she will assess whether what she uncovered is sufficiently similar to the applicant's innovation to justify a rejection of the patent application.

In general, two types of errors can be made: refusing a patent to a good innovation (type I error), or wrongly granting a patent to a non-patentable innovation (type II error). As a rejected patent can be re-applied for and will be scrutinized by another examiner, type II error is more common than type I error. Even though there is no consensus about the overall grant rate, many applications are eventually patented.[24] For this reason, we focus only on type II error.

The applicant gets a private benefit from his innovation whereas the examiner has social value in mind. Let $\overline{V}_G$ (resp., $\overline{W}_G$) be the private (resp., social) value of a good innovation that is granted a patent and $\overline{V}_R$ (resp., $\overline{W}_R$) when it is refused a patent, with $\overline{V}_G > \overline{V}_R \geq 0$ ($\overline{W}_G > \overline{W}_R$). The applicant derives higher benefit from a good patented innovation; but he can still get a positive benefit from his non-patented good innovation since it is new.

On the other hand, the private (resp., social) value of a bad innovation that is granted a patent is $\underline{V}_G$ (resp., $\underline{W}_G$) and 0 (resp., $\underline{W}_R$) if a patent is refused, with $\underline{V}_G > 0$ ($\underline{W}_R > \underline{W}_G$).[25] The applicant benefits from being granted a bad patent, as the probability of it being invalidated in court is relatively low, or a trial may invalidate only part of the claims.

The granting of a good patent has a higher (or potentially equal) social value than refusing a bad patent, which, in turn, has a higher value than refusing a good patent,

---

[21] In a survey of 4,269 examiners of the EPO reported by Friebel et al. (2006), 85.3 % of the respondents explained that high-quality work of applicants saved them a large amount of time in their work.

[22] Applications that contain a large amount of prior art might be more costly to process, and a different cost function could be used to capture this aspect of the process. Our findings depend on the functional form (1).

[23] Because examiners are different among themselves, they will interpret differently the information that they have. There is no objective criterion; and, as stated by Cockburn et al. (2003), there are "as many patent offices as patent examiners."

[24] By considering rejected applications that have been resubmitted, Quillen and Webster (2001, 2006) find that 87–95 % of U.S. applications are granted a patent, whereas Sampat and Lemley (2010) find a range of grant rates of 72–76 % for 2008.

[25] The upper (respectively, lower) bar indicates a good (respectively, bad) innovation, and the subscript G (respectively, R) indicates that the innovation is granted (respectively, refused) a patent.

which is also better than granting a bad patent: $\overline{W}_G \geq \underline{W}_R > \overline{W}_R > \underline{W}_G$. We define $\Delta W \equiv \underline{W}_R - \underline{W}_G > 0$, and we assume that $\Delta W < 1$, so as to obtain interior solutions for the optimal efforts.[26]

Private innovation values will likely be affected by cited prior art. The strength of a patent depends on whether prior art has been concealed to obtain the patent (Allison and Lemley 1998). The difference in value of two otherwise identical patents stems from the fact that it is easier for competitors to invalidate a patent on missing prior art than on cited prior art.

Uncited prior art can be due to a strategic concealment of information or to a lack of knowledge of it. In the former case, we explicitly introduce the cost associated with concealing prior art. To model *ex post* punishment for non-compliance, we assume that the private value of the patented innovation is negatively affected by strategic concealment. When prior art is concealed, the value of the patented innovation is $(1 - \alpha)V$, with $V = \{\overline{V}_G, \underline{V}_G\}$ and $\alpha \in [0, 1]$. The higher is α, the higher is the non-compliance punishment and, thus, the lower is the value of a patented innovation. We call the parameter α the "non-compliance punishment cost".

Prior art may also be missing because the applicant did not find it, and additional prior art might be revealed during a legal dispute. In the case of a good innovation, this additional prior art will not necessarily invalidate the patent; it will just make more costly a defense of a patent with less prior art. In the case of a bad patent, it will more likely be invalidated, since the missing prior art might be invalidating. We thus assume that an innovation that belongs to an abundant prior art field but patented with sparse prior art will be more costly to defend; it has a value $V'_G < V_G$ for $V_G = \{\overline{V}_G, \underline{V}_G\}$. We denote $\Delta \overline{V} = \overline{V}_G - \overline{V}'_G > 0$ and $\Delta \underline{V} = \underline{V}_G - \underline{V}'_G > 0$.

We summarize the values in the following matrix when a patent is granted with all prior art, with concealed prior art, or simply with missing prior art, and when a patent is refused.

	Grant a patent with all prior art	Grant a patent with concealed prior art	Grant a patent with missing prior art	Refuse a patent
Good innovation (p)	$\overline{V}_G, \overline{W}_G$	$(1 - \alpha)\overline{V}_G, \overline{W}_G$	$\overline{V}'_G, \overline{W}_G$	Not applicable No type I error
Bad innovation ($1 - p$)	$\underline{V}_G, \underline{W}_G$	$(1 - \alpha)\underline{V}_G, \underline{W}_G$	$\underline{V}'_G, \underline{W}_G$	$0, \underline{W}_R$

As applications differ in their informational content, so should the examiner's behavior. The PTO often claims that every application should receive identical scrutiny to guarantee that applicants are equally treated. Nevertheless, it seems

---

[26] The PTO is a "forward looking and long term" player that internalizes *ex ante* research incentives in the social value of a good innovation.

rational to treat two applications differently when both have the same social value but one is costlier to process than is the other.

We emphasize this issue by considering three different patenting processes: (*i*) the non-commitment case: the examiner only responds to the information transmitted; (*ii*) the second review process case: the examiner puts more emphasis on patents with little prior art content, as is the case, for instance, in emerging fields; and (*iii*) the commitment case: the examiner commits *ex ante* to specific levels of scrutiny depending on the prior art received.

To summarize, the timing of the game is as follows: In cases (*i*) and (*ii*).

- on date 1, the applicant decides how much effort $e \in [0, 1]$ to put into a prior art search. He finds prior art x and learns whether his innovation is patentable or not. He then files an application with announced prior art $\tilde{x}$;
- on date 2, the examiner observes the applicant's announcement $\tilde{x}$, and decides to undertake search effort $E \in [0, 1]$; and
- on date 3, depending on the examination outcome, the examiner grants a patent or not.

In case (*iii*).

- on date 1, the examiner commits *ex ante* to specific levels of scrutiny;
- on date 2, knowing these effort levels, the applicant chooses his search effort $e \in [0, 1]$. He finds x and learns whether his innovation is patentable or not. He files an application with announced prior art $\tilde{x}$; and
- on date 3, the examiner grants or refuses a patent based on the announced information, and the outcome of her own examination.

In the next two sections, we focus on the non-commitment case (case *i*). The second review process (case *ii*) and the commitment case (case *iii*) will be introduced in later sections. For future reference, we gather all of the parameters and variables that have been defined so far, and also those that will be introduced later in Table 1.

4 Technologies of Patent Examination and the Prior Art Report

We present the structure of the examination technology, and the applicant's prior art search technology when the examiner does not commit to any level of scrutiny efforts *ex ante* (case *i*).

4.1 Patent Examination Technology

After receiving an application with $\tilde{x}$, the examiner first updates her beliefs about its patentability, and chooses an effort level. Her updated beliefs depend crucially on her anticipation of the applicant's behavior. She draws some inference about patentability by considering the prior art that he has reported. If the applicant has an

Table 1 Definition of parameters and variables ($\widetilde{x}, e, E, \mu$ and θ are endogenous variables)

Parameter	Definition
p	Unconditional probability that an innovation is patentable or "good"
γ	Unconditional probability that the prior art is abundant
α	Non-compliance punishment cost
β	Second review screening
x	Prior art found by the applicant in his search $x \in \{x_l, 1\}$
$\overline{W}_G$	Social value that accrues from the grant of a patent to a good innovation
$\overline{W}_R$	Social value that accrues from the rejection of a good innovation
$\underline{W}_G$	Social value that accrues from the grant of a patent to a bad innovation
$\underline{W}_R$	Social value that accrues from the rejection of a bad innovation
$\Delta W \equiv \underline{W}_R - \underline{W}_G$	Incremental social value of the rejection of a bad innovation
$\overline{V}_G \ (\overline{V}_G')$	Private value of the grant of a patent to a good innovation (with missing prior art)
$\overline{V}_R$	Private value of the denial of a patent to a good innovation
$\underline{V}_G \ (\underline{V}_G')$	Private value of the denial of a patent to a bad innovation (with missing prior art)
$\Delta \overline{V} \equiv \overline{V}_G - \overline{V}_G'$	Incremental private value of a good patent (without missing prior art)
$\Delta \underline{V} \equiv \underline{V}_G - \underline{V}_G'$	Incremental private value of a bad patent (without missing prior art)
$\widetilde{x}$	Prior art that is released by the applicant in his application; $\widetilde{x} \in \{x_l, 1\}$
$\mu_h \ (\mu_l)$	Examiner's updated belief that the innovation is good when $\widetilde{x} = 1 (\widetilde{x} = x_l)$
e	Generic search effort of the applicant
e^t	Search effort of the applicant in a truthful revelation equilibrium
e^c	Search effort of the applicant in the commitment case
e_{SR}	Search effort of the applicant when the PTO implements a second review
E	Generic search effort of the examiner
$E_h \ (E_l)$	Search effort of the examiner when $\widetilde{x} = 1 (\widetilde{x} = x_l)$
$E_h^t \ (E_l^t)$	Search effort of the examiner in a truthful revelation equation when $\widetilde{x} = 1 \ (\widetilde{x} = x_l)$
$E_h^c \ (E_l^c)$	Committed search effort of the examiner when $\widetilde{x} = 1 (\widetilde{x} = x_l)$
$E_h^{SR} \ (E_l^{SR})$	Search effort of the examiner with a second review and when $\widetilde{x} = 1 \ (\widetilde{x} = x_l)$
$\theta \ (\theta^{SR})$	Probability that a bad applicant doesn't conceal some prior art (with a second review)

incentive to withhold some prior art, a (rational) examiner has no reason to believe that a given application has a probability p of being patentable.

Formally, μ_j for $j = l, h$ denotes the examiner's updated belief that the innovation is good, given that $\widetilde{x} = x_j$ has been revealed.[27] For any $\widetilde{x} \in X$ received, the examiner solves

$$\max_E [B(E) - C(E, \widetilde{x})],\tag{2}$$

where $C(E, \widetilde{x})$ is defined by (1) and

---

[27] The updating mechanism will be described later in Sect. 5.1., where μ_j is derived by using Bayes' rule.

$$B(E) = \mu_j \overline{W}_G + \left(1 - \mu_j\right)\left[E\underline{W}_R + (1 - E)\underline{W}_G\right]. \tag{3}$$

The first part of (3) represents the expected social value of a good innovation that generates a social value $\overline{W}_G$. The second part of (3) represents the expected social value of a bad innovation. With probability $\left(1 - \mu_j\right)$ the innovation is not patentable; the examiner finds it with probability E and refuses a patent, which generates a social value $\underline{W}_R$. With probability $(1 - E)$, she wrongly grants a patent to a bad innovation, which is worth $\underline{W}_G$ to society. As μ_j depends on $\widetilde{x}$, so does the gross benefit function (3).

The examiner's effort is E_h after receiving $\widetilde{x} = x_h \equiv 1$ or E_l after receiving $\widetilde{x} = x_l$. When the innovation is patentable, she can never prove otherwise. When it is not patentable, she can find invalidating prior art by exerting effort and refuse a patent. However, if she does not search enough, she may not find any evidence against patenting and will grant a patent to a bad innovation.

4.2 The Applicant's Prior Art Search and Report

The first applicant's decision is to search for prior art. As this is not our main focus, we analyze this decision in Sect. 5.2. We assume that $e > 0$, and we focus on the report of prior art, given that the applicant has already undertaken a search, found either x_l or 1, and learned whether his innovation is patentable. For expositional purposes, we call a "good" (resp., "bad") applicant an applicant who has learned that he has a patentable (resp., non-patentable) innovation.

If the applicant finds 1, he reports either x_l or 1. A good applicant always reports 1 as his expected gain from reporting 1 is $\overline{V}_G$ whereas it is $(1 - \alpha)\overline{V}_G < \overline{V}_G$ from reporting x_l. A bad applicant reports 1 only if $(1 - E_h)\underline{V}_G \geq (1 - E_l)(1 - \alpha)\underline{V}_G$ or, equivalently, if

$$(1 - E_h) \geq (1 - E_l)(1 - \alpha). \tag{4}$$

Inequality (4) represents the condition under which a bad applicant fully reports his prior art.

5 Non-commitment Case

We now characterize the applicant's optimal prior art revelation, and we derive the Perfect Bayesian equilibrium. Then, we determine the applicant's prior art search strategy.

5.1 Prior Art Revelation and Examination

We first consider that the applicant reports what he found, no matter what he learned in the search process. The examiner's updated beliefs, consistent with this behavior, are $\mu_h = \mu_l = p$. Therefore, conditional on receiving $\widetilde{x}$, the examiner's optimal efforts solution of (2) are

$$E_j^t = \widetilde{x}(1-p)\Delta W, \tag{5}$$

for $j = h, l$, $\widetilde{x} = \{x_l, 1\}$ and t stands for truthful. Note that 'truthful' here refers to truthful revelation of information: The applicant reveals all the prior art that he found. Even if he reports truthfully, he still might have a bad innovation, but did not find invalidating prior art.

We now define the equilibrium revelation strategy: A good applicant who finds 1 has no incentives to transmit less prior art than he found. A bad applicant who finds 1 reports $\widetilde{x} = 1$ if inequality (4) is satisfied or, by using (5), if $\alpha \geq \alpha_1$ where

$$\alpha_1 \equiv \frac{(1-p)(1-x_l)\Delta W}{1 - x_l(1-p)\Delta W}. \tag{6}$$

The threshold level α_1 is affected by changes in the parameters, as stated in the following Lemma (the proof of Lemma 1, as well as all the other proofs, is provided in the "Appendix"):

Lemma 1 *The threshold α_1 is decreasing with p and x_l and increasing with ΔW.*

As it is more likely that the innovation is good, the examiner has less incentives to search for invalidating information (from (5), E_l^t and E_h^t decrease), and it is more likely that applicants report truthfully. An increase in x_l increases E_l^t but has no effect on E_h^t. It follows that applicants report truthfully more often as α_1 is reduced. As the social value of rejecting a bad patent increases or the social value of wrongly granting a patent is reduced (i.e., ΔW increases), both efforts increase, which reduces the likelihood to report truthfully.

For high non-compliance punishment costs[28] ($\alpha \geq \alpha_1$), a bad applicant has too much to lose by concealing prior art; thus, an equilibrium exists in which applicants reveal truthfully. In fact, in a system where it is possible to enforce the doctrine of inequitable conduct, applicants would always report all the prior art. We summarize these findings in the following proposition:

Proposition 1 *For high non-compliance costs ($\alpha \geq \alpha_1$), an equilibrium exists in which the applicant always reveals the prior art that he found (x_l or 1). The examiner exerts a higher scrutiny effort when she receives more prior art ($E_h^t > E_l^t$).*

In an equilibrium in which bad applicants report truthfully, the examiner intensifies her scrutiny effort when she receives more prior art ($E_h^t > E_l^t$), as it is less costly to do so. This result holds when $\alpha \geq \alpha_1$. However, the non-compliance cost α is likely to be quite small as it is difficult for the PTO to establish that some prior art has been concealed, and the PTO does not enforce the doctrine of inequitable conduct (Kesan 2002).

For $\alpha \in [0, \alpha_1[$, a bad applicant realizes that concealing prior art is best for him; inequality (4) is violated, and full disclosure cannot be an equilibrium. However, if full concealment is adopted, all applications with 1 are patentable. A rational examiner adjusts her behavior by exerting no effort on these applications and focusing solely on applications with x_l. But then, a bad applicant realizes that he can

---

[28] In terms of reputation, or direct penalties in the case of a lawsuit.

obtain a patent for sure by submitting 1. Hence, full concealment cannot be an equilibrium either. We summarize these findings in the following Lemma:[29]

Lemma 2 *For low non-compliance costs* $(0 \leq \alpha < \alpha_1)$, *no equilibrium in pure strategies exists.*

When it is not costly to conceal prior art, applicants have incentives sometimes to not report all of the prior art, but they may sometimes report truthfully. Therefore, we consider a Perfect Bayesian Equilibrium in which a bad applicant randomizes his report decision. Let θ be the probability that a bad applicant who finds 1 reports $\widetilde{x} = 1$, and $(1 - \theta)$ the probability that he conceals some prior art and reports $\widetilde{x} = x_l$.

When the examiner receives an application containing $\widetilde{x} = 1$, it can originate from a good or a bad applicant. When the application contains $\widetilde{x} = x_l$, it can come from a (good or bad) applicant who has discovered x_l, but also from a bad applicant who has found 1 but revealed $\widetilde{x} = x_l$. The examiner's updated beliefs consistent with randomization are[30]

$$
\begin{aligned}
\mu_h(\theta) &= prob(good \mid \widetilde{x} = 1) = \frac{p}{p + \theta(1-p)} > p, \\
\mu_l(\theta) &= prob(good \mid \widetilde{x} = x_l) = \frac{p(1 - e\gamma)}{1 - \gamma e(p + \theta(1-p))} < p.
\end{aligned}
\tag{7}
$$

In this equilibrium, a bad applicant is indifferent between revealing $\widetilde{x} = 1$ and getting $(1 - E_h(\theta))\underline{V}_G$, and revealing $\widetilde{x} = x_l$ and getting $(1 - E_l(\theta))(1 - \alpha)\underline{V}_G$. The efforts $E_j(\theta)$ for $j = l, h$ are the solution of (2) with updated beliefs (7). A unique value of $\theta \in [0, 1]$ exists that satisfies

$$
1 - E_h(\theta) = (1 - E_l(\theta))(1 - \alpha).
\tag{8}
$$

We posit the following result:

Lemma 3 *For low non-compliance costs* $(0 \leq \alpha < \alpha_1)$, *there exists a unique probability* θ^* *with which a bad applicant who finds 1 discloses it. This probability increases with* α, γ, *and* e.

This result is graphically illustrated in Fig. 1, where $E_h(\theta)$ and $E_l(\theta)$ are functions of θ. We represent θ^* as the solution of (8) for the optimal values E_h^* and E_l^*.

The probability θ^* measures the applicant's *level of candor*. An increase in α has a direct effect on θ: Fewer bad applicants report $\widetilde{x} = x_l$ when they find 1. When it becomes too costly to report strategically, then $\theta^* = 1$, which corresponds to an equilibrium of truthful revelation.

---

[29] The non-existence of an equilibrium in pure strategies is akin to an audit without commitment (Khalil 1997).

[30] These beliefs are calculated for any given e. However, e denotes the anticipated search effort of the applicant which (in equilibrium) will be the equilibrium search effort.

Fig. 1 Optimal level of θ^*
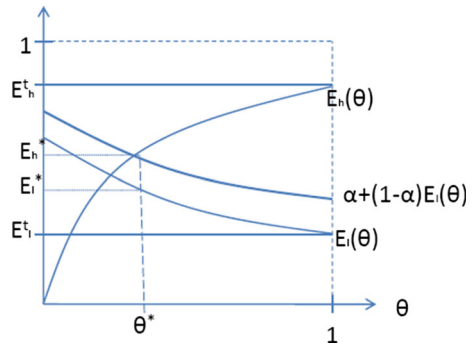
As it is more likely that the innovation belongs to a rich prior art field, for a given θ, the function $E_h(\theta)$ is not affected (as $\mu_h(\theta)$ does not depend on γ) while $E_l(\theta)$ increases. However, anticipating that there will be more scrutiny when he reveals x_l, a bad applicant who found 1 reveals more often 1 and thus θ^* increases.

As e increases, the applicant is more likely to find 1; thus, the examiner's updated belief to have a good innovation with 1 decreases. This pushes the examiner to increase E_l: which (in turn) pushes a bad applicant to report 1 when he finds 1 and thus θ^* increases.

We present the equilibrium in the following proposition:

Proposition 2 *For low non-compliance costs $(0 \leq \alpha < \alpha_1)$, a semi-separating equilibrium exists in which a bad applicant randomizes his prior art revelation with probability θ^*, whereas a good applicant always reports all his prior art. The examiner exerts a higher search effort when she receives more prior art: $E_h^* \geq E_l^*$.*

This finding emphasizes the main strategic tension that faces an applicant who has discovered that his innovation is non-patentable.[31] It is consistent with empirical findings by Lampe (2012): Applicants do withhold citations to relevant prior art.[32] We compare the efforts and perform some comparative statics as reported in the following corollary:

---

[31] To simplify, we restrict reports to be either $\widetilde{x} = 1$ or $\widetilde{x} = x_l$. Thus, an applicant who found 1 can report either $\widetilde{x} = 1$ or $\widetilde{x} = x_l$, while an applicant who found x_l can only report $\widetilde{x} = x_l$. In a relaxed framework where an applicant who found 1 can report any value $\widetilde{x} \leq 1$, a good applicant will report $\widetilde{x} = 1$ as under-reporting prior art (when an innovation is good) hurts the innovator in subsequent litigation. Unlike the good applicant, a bad applicant who found 1 will still have an incentive to withhold prior art to have a better chance to obtain a patent. However, reporting a different level than x_l cannot be an equilibrium since it then fully reveals his type. Hence, mimicking the applicant who reports $\widetilde{x} = x_l$ would still be an equilibrium.

[32] We can also interpret α_1 as a measure of the likelihood of strategic behavior, where x_l could represent the importance of missing prior art in helping to determine the innovation's novelty (information cues). A smaller x_l indicates that missing prior art is becoming more useful to the examiner in assessing the patentability, and applicants take advantage of this by withholding prior art more often (Lemma 1). When concealing prior art can strongly impact an examiner's screening, applicants are more likely to withhold it for strategic reasons.

Corollary 1 *The optimal effort levels of the examiner are such that* $E_h^t \geq E_h^* \geq E_l^* \geq E_l^t$ *and*

- *when the probability of having a good innovation p increases,* E_h^* *and* E_l^* *decrease;*
- *when the probability of having an innovation in a rich prior art field* γ *increases,* E_h^* *and* E_l^* *increase;*
- *when the non-compliance cost* α *increases,* E_h^* *increases while* E_l^* *decreases; and*
- *when the effort of the applicant e increases,* E_h^* *and* E_l^* *increase.*

These findings illustrate the mechanism of the examiner's effort allocation. She still exerts a higher effort when receiving 1 rather than x_l ($E_h^* \geq E_l^*$), as it is less costly to do so; but now she intensifies her effort when receiving x_l compared to the truthful revelation equilibrium effort ($E_l^* \geq E_l^t$) as more bad applicants are likely to report x_l. At the same time, she also exerts less effort when she receives 1 ($E_h^t \geq E_h^*$) as these applications are more likely to be good.

The probability of having a patentable innovation p plays a central role in the examiner's provision of effort. When p increases, less bad applicants will apply for a patent; thus, the examiner finds it less attractive to exert effort irrespective of the amount of prior art that has been submitted.[33]

When γ increases, there are fewer applications with x_l. While good applications contain abundant prior art, only a fraction θ of the bad ones have abundant prior art. As such, the probability of a good application with sparse prior art, μ_l, decreases, and E_l^* increases. Furthermore, γ influences E_h^* through θ, which increases when γ increases; E_h^* increases as well.

When the non-compliance cost α increases, strategic behavior is naturally deterred (θ^* increases). The probability of having a good application with abundant prior art, μ_h, decreases. The examiner, thus, increases her effort E_h^* to screen these applications. Conversely, the probability of having a good application with x_l increases; thus, her effort E_l^* decreases.

Finally, as the applicant's effort e increases, the updated belief of having a good innovation when receiving x_l decreases; thus, the examiner intensifies her search E_l. However, θ^* increases as well; thus, less strategic concealment occurs. This pushes the examiner to intensify her search when receiving 1 as well.

5.2 Applicant's Search of Prior Art

The scarcity of prior art reported should not be taken as evidence that applicants do not search for it. We now determine the applicant's effort to search for prior art by focusing only on the case of strategic report of prior art (i.e., $0 \leq \alpha < \alpha_1$). The applicant's benefit function is

---

[33] In a previous version of the paper (available upon request), we allow the examiner to reject good applications. When type I errors are permitted, the examiner's screening effort could increase with the quality of applications.

$$\Pi(e) = e\gamma[p\overline{V}_G + (1-p)((1-E_h)\theta + (1-\theta)(1-E_l)(1-\alpha))\underline{V}_G] + (1$$
$$- e)\gamma[p\overline{V}'_G + (1-p)(1-E_l)\underline{V}'_G] + (1-\gamma)[p\overline{V}_G + (1-p)(1-E_l)\underline{V}_G]$$
$$- c(e).$$

His effort e generates a probability of finding prior art conditional on its existence. Thus, after making an effort e, with probability e the applicant finds 1 (resp., x_l) if the prior art is abundant (resp., scarce) with probability γ (resp., $1 - \gamma$).[34] He then files for a patent with some prior art.

In the case of abundant prior art, if he discovers that his innovation is good (with probability p), he will apply for a patent with 1, and the patent will be granted. If his innovation is bad (with probability $1 - p$), he will randomize his prior art report to increase his chances of being granted a patent: with probability θ (resp., $1 - \theta$) he reports 1 (resp., x_l). The examiner can wrongly grant a bad patent. In equilibrium, this behavior is anticipated by the examiner, who adjusts her beliefs such that (8) holds.

In the second part of the benefit function, the applicant finds x_l with probability $(1 - e)$ in an abundant prior art field (with probability γ). He is granted a deserving patent with probability p or is wrongly granted a patent with probability $(1 - p)(1 - E_l)$. In the third part of the benefit function, the innovation belongs to a scarce prior art field (with probability $1 - \gamma$), and the applicant can be rightly (resp., wrongly) granted a patent with probability p (resp., $(1 - p)(1 - E_l)$). Whenever there is (unintentional) missing prior art, the private value of the patent is V'_G.

By using (8), and given the examiner's rational anticipations, the applicant's maximization program simplifies to

$$\max_e \Pi(e, e^*) = e\gamma[p\overline{V}_G + (1-p)(1-E_h(e^*))\underline{V}_G] + (1-e)\gamma$$
$$\times \left[p\overline{V}'_G + \frac{1-p}{1-\alpha}(1-E_h(e^*))\underline{V}'_G\right](1-\gamma)\left[p\overline{V}_G + \frac{1-p}{1-\alpha}(1-E_h(e^*))\underline{V}_G\right] - c(e),$$
$$(9)$$

and the first order condition is

$$\gamma p\Delta\overline{V} + (1-p)\gamma(1-E_h(e^*))(\underline{V}_G - \frac{\underline{V}'_G}{1-\alpha}) = c'(e). \qquad (10)$$

We assume that the examiner has rational anticipations regarding θ^* and the updated beliefs (7). The applicant cannot influence E_h, which depends on the equilibrium effort's level e^*, which is anticipated by the examiner. Nevertheless, through his effort e, he can influence the probability to find more prior art. He is inclined to search if it uncovers (good) patent applications with abundant prior art (first term of (10)). But, he also has an anti-social goal in mind, since a bad innovation can still be granted a patent with abundant prior art (second term of (10)).[35]

---

[34] In particular, note that if $e = 0$, the applicant will always find $x_l > 0$. This assumption seems realistic since the applicant must have some prior art knowledge regarding what his innovation is improving upon.

[35] However, when α is large, the incentive to search in order to patent bad innovations will disappear.

Proposition 3 *For low non-compliance costs* $(0 \leq \alpha < \alpha_1)$, *a unique strictly positive equilibrium effort* e^* *(decreasing with* α*) exists, if*

$$p\Delta\overline{V} + (1-p)[\Delta\underline{V} - (1-p)\Delta W(\underline{V}_G - x_l\underline{V}'_G)] > 0. \tag{11}$$

The applicant exerts a positive effort if the gain from being granted a bad patent while reporting 1 is much higher than the gain from being granted a bad patent with missing prior art (i.e., if $\underline{V}_G$ is much higher than $\underline{V}'_G$). *Ex ante*, the applicant has an incentive to search if he gets some benefit even in the case of a bad innovation. When the cost of concealing prior art α increases, the incentive to search is weaker.

If (11) is violated (if x_l or $\Delta\underline{V}$ is small), a threshold $\widetilde{\alpha}$ exists, with $0 < \widetilde{\alpha} < \alpha_1$ such that a positive effort will still be induced if $\alpha < \widetilde{\alpha}$. On the other hand, if $\alpha \geq \widetilde{\alpha}$, the examiner will choose to reduce her search effort when receiving x_l (as the cost of effort becomes very large). In turn, this provides incentives to the applicant not to search and report x_l. In other words, the applicant stops concealing prior art simply because he does not search for it in the first place.

In the next two sections, we study two mechanisms that are designed to promote the flow of information from applicants to the examiner.

6 A Second-Review Process

In March 2000, in reaction to numerous quality-related criticisms about the granting of business-method patents (main class 705), the PTO began a quality patent improvement initiative, called the "Second Pair of Eyes Review" (SPER). It involves several measures, such as the hiring of better-trained examiners, an obligation to consult non-patent prior art sources and, more importantly, a second-level of examination of patents that were granted within the main 705 classification.

To study the effects of this initiative empirically, Allison and Hunter (2006) analyze the composition of applications before and after the SPER initiative. They argue that an applicant endowed with a business-method innovation can choose to submit his application in the main class 705 or in other main classes that are related to business methods (with a secondary 705 classification).

In a variant of our model that incorporates a second review process similar to the SPER, we investigate its impact on the patenting process and the behavior of applicants. We model this initiative with a second review of awarded patents when x_l has been revealed. While the PTO could, at a very substantial cost, subject all awarded patents to a second review, this mechanism allows her to 'target' applications with sparse prior art.

Innovations can belong to two technologically related fields: In one of them, there is less prior art to be found. By strategically drafting their applications, applicants can "opt out" of the field where there is abundant prior art. We assume that when patents are granted to bad innovations the second review eliminates a fraction $\beta \in [0, 1]$ of them. If a patent is granted to a good innovation, the second review will not invalidate it. We call β the second-review screening.

Similarly to the non-commitment case, an equilibrium in which the applicant reports truthfully exists, so long as $(1 - E_h) \geq (1 - E_l)(1 - \beta)(1 - \alpha)$ is satisfied [(inequality similar to (4)] or, equivalently, for $\alpha \geq \alpha_1(\beta)$ with

$$\alpha_1(\beta) \equiv \frac{(1 - p)(1 - (1 - \beta)x_l)\Delta W - \beta}{(1 - \beta)(1 - x_l(1 - p)\Delta W)} \leq \alpha_1.$$

For $\alpha < \alpha_1(\beta)$, the probability θ is the solution of

$$1 - E_h(\theta) = (1 - E_l(\theta))(1 - \beta)(1 - \alpha). \tag{12}$$

We show that there exists a unique θ_{SR}^* such that (12) is satisfied. The optimal efforts E_h^{SR} and E_l^{SR} are the solution of the program (2). We illustrate graphically these results by representing $E_h(\theta)$ and $E_l(\theta)$ as a function of θ in Fig. 2, and we plot the optimal efforts E_h^{SR} and E_l^{SR}.

Although β can take any value between 0 and 1, any $\beta \geq \bar{\beta}$ where $\bar{\beta} = (E_h - E_l - \alpha(1 - E_l))/(1 - E_l)(1 - \alpha)$ induces $\theta_{SR}^* = 1$ (no strategic concealment), whereas for $\beta \in [0, \bar{\beta}]$, $\theta_{SR}^* \in [\theta^*, 1]$. As $\theta_{SR}^* \geq \theta^*$, in the presence of a second review, the applicant behaves strategically less often. Furthermore, θ_{SR}^* increases with β, so that an increase in the precision in the second review reduces the strategic concealment of information.

A comparison of the efforts with and without a second review leads to the following findings:

Proposition 4 *For low non-compliance costs* $(0 \leq \alpha < \alpha_1(\beta))$, *the introduction of a second review reduces the strategic concealment of prior art* $(\theta_{SR}^* \geq \theta^*)$. *The examiner increases her scrutiny effort when she receives 1* $(E_h^{SR} > E_h^*)$, *and reduces it when she receives* x_l $(E_l^{SR} < E_l^*)$.

With a second review, applications with x_l get more scrutiny; thus, the applicant has less incentive to conceal prior art, and he will reveal it more often $(\theta_{SR}^* \geq \theta^*)$. Interestingly, it changes the examiner's behavior as well. As she is more likely to receive bad applications with 1, the examiner intensifies her scrutiny of these
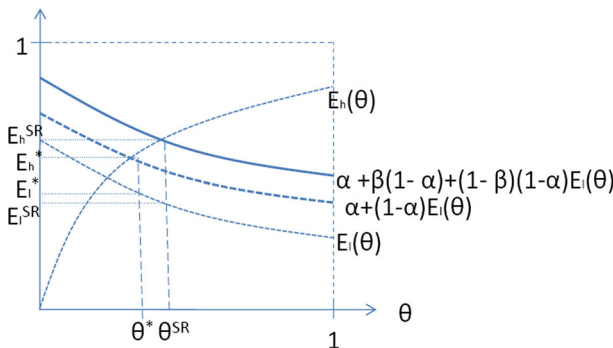

Fig. 2 SPER

applications ($E_h^{SR} > E_h^*$), while she reduces her scrutiny of applications with x_l that are now more likely to be good ($E_l^{SR} < E_l^*$).

Less concealment occurs when a second-review is implemented. Indeed, for $\alpha \in [\alpha_1(\beta), \alpha_1[$, an equilibrium exists in which the applicant reveals truthfully under a second-review regime but would conceal information in absence of a second review. For $\alpha \in [0, \alpha_1(\beta)[$ the applicant conceals prior art under both regimes; but compared to the non-commitment case, his level of candor θ_{SR}^* is increased when the second review is implemented.

For the sake of simplicity, we ignore the cost of implementing the second review policy. It is clear that the introduction of a second examination is costly for the PTO. The introduction of such a cost would make our model intractable without changing the flavor of the results. Our findings would be more dramatic, as it would make the screening of patents with abundant prior art costly and, therefore, would reduce the examiner's efforts when prior art is abundant.

While this policy reduces strategic concealment, it also induces the applicant to provide less effort, as is stated in the following proposition:

Proposition 5 *Assume that $\alpha = 0$. The introduction of a second review reduces the effort of the applicant ($e_{SR}^* < e^*$).*

With a second review, the applicant has less incentives to search for prior art. No matter what he reports, a bad applicant who finds 1 is more likely to be refused a patent. This is because, compared to a single review, a second review of applications with x_l also increases scrutiny when 1 is reported; the applicant has less incentives to search. Thus, by reducing his search effort, he is less likely to find 1, and he reports x_l more often. He might be granted a patent by mistake, as the examiner search intensity is reduced (even if there is a second review).

To analyze the effects of a second review, we compare the number of applications and granted patents, as well as a proxy of their quality before and after its implementation (calculations are provided in the "Appendix"). Our findings are summarized in the following lemma:

Lemma 4 *For low non-compliance costs ($0 \le \alpha < \alpha_1(\beta)$), for any $\beta > 0$, the implementation of the second-review initiative has the following effects*:

- *Fewer (resp., more) applications with low (resp., high) prior art content are scrutinized;*
- *Fewer patents with abundant prior art are granted; and*
- *The ratio of good patented innovations over all patented innovations decreases (resp., increases) for applications with abundant prior art (resp., scarce prior art).*

As less concealment of information occurs with a second review, more (bad) applications are filed with abundant prior art. Even though there is more scrutiny, overall since there are more bad applications, fewer good patents are granted in an abundant field. This lower ratio of good patents over all patents with prior art 1 is

clearly one unintended effect of the higher scrutiny of applications that report x_l. As such, one contribution of our analysis is to show that the merit of such a reform should not be judged only by its (positive) direct impact on patents with scarce prior art, but also by its impact on the quality of awarded patents with abundant prior art.

7 An *Ex Ante* Commitment by the PTO

In the previous sections, we have assumed that the examiner's scrutiny level is an optimal response to every application that she receives. While this may be realistic in many cases, it is also interesting to investigate the case where the examiner commits to specific scrutiny levels, potentially contingent on the prior art reported. We analyze an *ex ante* commitment policy in which the applicants' search effort and concealment strategies become responses to this commitment.

Let E_h^c and E_l^c be the examiner's effort levels when, respectively, $\widetilde{x} = 1$ and $\widetilde{x} = x_l$ are reported by the applicant; and, to simplify, we assume that concealment is costless ($\alpha = 0$). Two qualitatively different cases must be considered: either $E_l^c \geq E_h^c$ or $E_l^c < E_h^c$.

We first consider that $E_l^c < E_h^c$. If the applicant finds 1, he learns whether his innovation is patentable or not. A good applicant will report all of the prior art, whereas a bad applicant may decide to transmit only x_l. The next result is the first step in finding the optimal policy:

Lemma 5 *Consider a policy* $\left(E_l^c, E_h^c\right)$ *such that* $E_l^c < E_h^c$, *then it is always possible to find a policy with strictly lower search costs.*

A policy involving $E_l^c < E_h^c$ is cost-dominated by another policy with a smaller gap between E_h^c and E_l^c (see "Appendix"). Intuitively, reducing E_h^c does not change the optimal reporting strategy (all bad applicants conceal prior art), but it reduces scrutiny costs.[36]

The commitment policy that involves $E_l^c \geq E_h^c$ induces truthful revelation of information and the underlying commitment problem is less intense. Indeed, in a truthtelling equilibrium, $\left(E_l^c, E_h^c\right)$ are *ex post* inefficient; but the gains from reneging (switching to E_h^t and E_l^t) are lower. Thus, compared to $E_l^c < E_h^c$, the examiner faces a low-intensity commitment problem when $E_l^c \geq E_h^c$.

When $E_h^c = E_l^c$, the applicant is indifferent between reporting truthfully $\widetilde{x} = 1$ and strategically $\widetilde{x} = x_l$. When he is indifferent, he reports truthfully and the examiner's gross benefit $B(E_j^c)$ for $j = h, l$ is defined by Eq. (3) for $\mu_j = p$.

The commitment policy that induces search and truthful revelation is the solution of

$$\max_{E_h^c, E_l^c} e\left(E_h^c, E_l^c\right) \gamma \left[B\left(E_h^c\right) - C\left(E_h^c, x_h\right)\right] + \left(1 - e(E_h^c, E_l^c)\gamma\right)\left[B\left(E_l^c\right) - C\left(E_l^c, x_l\right)\right]$$

subject to $E_h^c \leq E_l^c$, where $C\left(E, x_j\right)$ for $j = h, l$ is defined by (1).

---

[36] Another drawback is that the examiner's commitment problem is very intense. Indeed, the examiner would gain $C(E_h^c)$ by reneging on her commitment, without altering the screening intensity.

The next result simplifies the analysis of the examiner's program:

Proposition 6 *If the examiner wants to induce truthful reports, the unique optimal commitment policy is such that $E_h^c = E_l^c \equiv E^c$.*

For an applicant who has discovered $x = 1$ but has learned that his innovation is not patentable, E_l^c and E_h^c represent, respectively, the probabilities with which the non-patentability of his innovation will be discovered if he conceals prior art and if he reports truthfully. If $E_l^c < E_h^c$, the applicant is naturally inclined to withhold prior art and to report x_l as it results in a lower probability of patent rejection. However, when E_l^c approaches E_h^c, the incentives to conceal prior art are reduced; and, if $E_l^c = E_h^c$ he will be indifferent between revealing prior art or not (for definiteness we assume he reveals).

It is intuitive that committing to levels such that $E_l^c > E_h^c$ will continue to elicit truthful reports. However these levels are at odds with the *ex post* efficient ones (determined in Proposition 2) that are such that $E_l^* < E_h^*$. Combining these two inequalities, in the commitment case, effort levels that elicit truthful reports and that involve the least *ex post* inefficiencies must thus verify $E_l^c = E_h^c$. By using this result, the program of the examiner becomes

$$\max_{E^c} \left[B(E^c) - \frac{(E^c)^2}{2x_l} + e(E^c)\gamma \frac{1 - x_l (E^c)^2}{x_l} \frac{1}{2} \right]. \tag{13}$$

The first term represents the (gross) social benefit of screening an application. The second term is the baseline cost that is incurred when an application with x_l is submitted. The last term represents the cost savings for the PTO when an applicant searches and discloses all of the prior art.

Unlike in the non-commitment case, the applicant's search effort depends on E^c. He chooses $e(E^c)$ as the solution of the maximization of $\Pi^c(e)$ (defined in the "Appendix"), which is

$$\gamma p \Delta \overline{V} + (1 - E^c)(1 - p)\gamma \Delta \underline{V} = c'(e). \tag{14}$$

We denote e^c the solution of (14), and we differentiate it with respect to E^c which yields

$$\frac{de^c}{dE^c} = -(1 - p)\gamma \frac{\Delta V}{c''(e^c)} < 0. \tag{15}$$

Expression (15) shows the impact of a variation in the examiner's scrutiny on the applicant's search effort: A stronger commitment to decrease her scrutiny increases the applicant's incentives to search for prior art. Computing the first-order condition of (13) and using (14) we obtain

$$(1 - p)x_l \Delta W = E^c - \gamma(1 - x_l)E^c \left(e^c - (1 - p)\frac{\Delta V}{c''(e^c)} \frac{E^c}{2} \right). \tag{16}$$

A comparison of (16) and (5) shows that, unlike in the non-commitment case, the examiner has incentive motives when she chooses her scrutiny level. As evidenced

by (15), by committing to reduce her effort, the examiner boosts the applicant's incentive to search for prior art and reduces her search cost. However, it also results in a larger number of bad patents. The following result characterizes the examiner equilibrium scrutiny:

Proposition 7 *The examiner's program* (13) *has a unique interior solution. Moreover, the equilibrium scrutiny level* E^c *is increasing in* ΔW *and* $\Delta \overline{V}$.

When the social value of an innovation (ΔW) increases, the examiner increases her scrutiny because there is a direct benefit to do so. This is similar to the non-commitment case. While in the non-commitment case, an increase in γ increases scrutiny effort levels, such an increase has an ambiguous effect in the commitment case where prior art concealment no longer arises.

The applicant's incentives to search for prior art mostly depend on the increment in value that is brought by patents with abundant prior art: ΔV. The following proposition summarizes the equilibrium efforts and compares them to those obtained in the non-commitment case:

Proposition 8 *In the commitment case, for* $\alpha = 0$ *(non-compliance cost is null), there exists two distinct levels of* $\Delta \underline{V}$, *denoted by* $\underline{\Delta}$ *and* $\overline{\Delta}$, *such that*

- *for* $\Delta \underline{V} > \overline{\Delta}$, *the examiner reduces her effort,* $E_h^* > E_l^* > E^c$, *and the applicant intensifies his effort,* $e^c > e^*$; *and*
- *for* $\Delta \underline{V} < \underline{\Delta}$, *the examiner increases her effort,* $E^c > E_h^* > E_l^*$, *and the applicant reduces his effort,* $e^c < e^*$.

The main trade-off of the PTO is between intense prior art searching and committing to strong audit standards. Whether the PTO chooses to favor a particular strategy depends on the incremental *private value* of finding prior art for a non-patentable innovation. When this value is very high ($\Delta \underline{V} > \overline{\Delta}$), the applicant has strong incentives to discover prior art; and, as is shown by Eq. (15), by lowering her effort E^c the examiner induces the applicant to increase his search effort e^c. This type of commitment is akin to a *delegation* model where the principal commits to interfere less in order to boost subordinate initiative (Aghion and Tirole 1997).[37]

When $\Delta \underline{V} < \overline{\Delta}$, (15) shows that reducing E^c will have a modest effect on the applicant's search effort. As such, unlike in the previous case, the examiner prefers to commit to a strong level of patent screening. In this case, $E^c > E_h^* > E_l^*$; and the objective of the PTO is to screen out bad applications.

The PTO's optimal commitment can be discussed along other parameters as well. For instance, if the applicant's marginal cost of prior art search $c''(e^c)$ is small, the examiner will be more likely to delegate search to the applicant by lowering her audit effort level E^c.

Whether $\Delta \underline{V}$ is small or large remains an empirical question, and it depends on the likelihood of patent invalidation when prior art is missing.

---

[37] However Aghion and Tirole (1997)'s model does not feature information concealment.

8 Conclusion

A patent examiner must assess an innovation's patentability by studying related information that is already in the public domain, which is termed "prior art". When the application contains little prior art and when the prior art is sparse or is not easily accessible, it may be difficult to judge the novel content of the innovation.

Prior art may be difficult to gather for examiners. In particular, if most of it is non-patent information, examiners may have difficulty in assessing the patentability of the innovation. This will lead to a reduction in patent quality, as more questionable patents are issued. There needs to be incentives for applicants to provide relevant information to help examiners make accurate decisions.

We analyze issues that are related to the gathering and revelation of prior art. The applicant's decision to reveal prior art depends on the information that is gathered and the patent screening policy that is adopted by the PTO. After receiving an application, the examiner undertakes a costly search effort that depends on her beliefs with regard to patentability, and decides whether to grant a patent.

We consider different mechanisms: one in which the PTO cannot commit to any search effort; one in which there is a second review for some applications; and one in which the examiner commits *ex ante* to a screening intensity.

In the first case, a bad applicant might conceal prior art to increase his chances of obtaining a patent. The examiner intensifies her search when she receives more prior art. Thus, she devotes more effort to applications that, more likely, come from good applicants.

In the second case, the introduction of a more thorough screening when little prior art is revealed reduces information concealment, but its overall success remains unclear. Indeed, the imposition of stricter screening when prior art is scarce may improve the quality of patents in this field, but could negatively affect the quality in another field.

In the third case, by committing *ex ante* to the same screening effort, the examiner induces the truthful revelation of information. A commitment to a lower level of screening may increase the amount of prior art that is transmitted to the PTO.

We make the simplifying assumption that the PTO and examiners have the same objective functions and, thus, that they represent the same decision maker. However, this is unlikely as the PTO is a federal agency that employs about 9,300 patent examiners (as of 2014) who get rewards for reaching predetermined goals. Therefore, examiners' search efforts are probably influenced by their incentive schemes. They may also have alternative objectives, such as career prospects. These issues are the focus of future research.

Appendix

Proof of Lemma 1 If the applicant reports truthfully, from (5), $E_h^t > E_l^t$ as $x_l < 1$. Inequality (4) can be rewritten as $\alpha \geq \alpha_1 = (E_h^t - E_l^t)/(1 - E_l^t)$. By replacing E_h^t and E_l^t, α_1 becomes (6). The derivatives of α_1 with respect to the different parameters are $\partial \alpha_1/\partial p < 0$, $\partial \alpha_1/\partial x_l < 0$, and $\partial \alpha_1/\partial \Delta W > 0$.

Proof of Proposition 1 Following the proof of Lemma 1, as long as $\alpha \geq \alpha_1$, $E_h^t > E_l^t$. Furthermore, $\partial E_j^t/\partial p < 0$, $\partial E_j^t/\partial \Delta W > 0$ for $j = h, l$ and $\partial E_l^t/\partial x_l > 0$.

Proof of Lemma 2 No equilibrium in pure strategies for $\alpha < \alpha_1$ Another possible equilibrium is one in which the applicant always conceals the prior art. The updated beliefs of an examiner who believes that a bad applicant who finds 1 always conceals it are $\mu_h = p\gamma/p\gamma = 1$ and $\mu_l = p(1 - \gamma)/(1 - \gamma p) < p$. If she receives 1, the examiner prefers not to search, as only a good applicant will reveal 1. If she receives x_l, $E_l = x_l(1 - \mu_l)\Delta W < 1$.

A good applicant has no incentive to deviate while a bad applicant is always expected to conceal prior art. Thus, we should have $(1 - E_l)(1 - \alpha) > (1 - E_h)$. However, as $E_h^* = 0$, it is never satisfied. Therefore, the applicant may decide to report 1 to fool the examiner who believes that only good applicants report 1. If he reports 1, he gets a patent as the examiner believes he has a good innovation and reports truthfully. This confirms that this is not an equilibrium.

By using the same argument, we show that the examiner will not make any effort when she observes 1, but she also will not award a patent. When she observes x_l, she makes a positive effort. A good applicant will not deviate, whereas a bad one will deviate, as he can get a patent by fooling the examiner. This is not an equilibrium.

The examiner may also believe that no applicants will report truthfully. By using the same argument we show that this cannot be an equilibrium.

Proof of Lemma 3 and Proposition 2: Subgame Perfect Bayesian Equilibrium We prove the existence of a semi-separating equilibrium in which a bad applicant randomizes his revelation decision. In this case (7) and (8) must be satisfied, $E_h(\theta) = (1 - \mu_h(\theta))\Delta W$, and $E_l(\theta) = x_l(1 - \mu_l(\theta))\Delta W$, where $\mu_h(\theta)$ and $\mu_l(\theta)$ are defined by (7).

We calculate that $\partial \mu_h(\theta)/\partial \theta < 0$, $\partial^2 \mu_h(\theta)/\partial \theta^2 > 0$, $\partial \mu_l(\theta)/\partial \theta > 0$, and $\partial^2 \mu_l(\theta)/\partial \theta^2 > 0$. Therefore, $\partial E_h(\theta)/\partial \theta > 0$, $\partial^2 E_h(\theta)/\partial \theta^2 < 0$, $\partial E_l(\theta)/\partial \theta < 0$, and $\partial^2 E_l(\theta)/\partial \theta^2 < 0$.

When $\theta = 0$, $E_h(0) = 0 < E_l(0) = x_l(1 - \mu_l(0))\Delta W$; and when $\theta = 1$, $E_h(1) = E_h^t > E_l(1) = E_l^t$. From equality (8), θ must satisfy $E_h(\theta) = \alpha + (1 - \alpha)E_l(\theta)$. As $E_h(\theta)$ is strictly increasing and convex, and $E_l(\theta)$ is strictly decreasing and convex, there exists a unique θ^*. We illustrate this finding in Fig. 1. We verify that when $\theta = \theta^*$, $E_h(\theta^*) = \alpha + (1 - \alpha)E_l(\theta^*) > E_l(\theta^*)$.

Proof of Corollary 1 It is deduced from the previous proof, as $E_h(0) = 0 < E_h(1) = E_h^t$ and $E_l(0) > E_l(1) = E_l^t$. Therefore, for any θ, $E_h(\theta) \leq E_h^t$ and $E_l(\theta) \geq E_l^t$. Thus $E_h^t \geq E_h^* \geq E_l^* \geq E_l^t$. Static comparatives are summarized in Table 2.

Table 2 Comparative statics

parameter y	$\frac{\partial E_l(\theta)}{\partial y}$	$\frac{\partial(\alpha+(1-\alpha)E_l(\theta))}{\partial y}$	$\frac{\partial E_h(\theta)}{\partial y}$	$\frac{\partial \theta^*}{\partial y}$
p	−	−	−	?
γ	+	+	0	+
α	0	+	0	+
ΔW	+	+	+	?
e	+	+	0	+

Proof of Proposition 3   The applicant's benefit, $\Pi(e)$, is a function of e and of the examiner's (rational) expectation e^*. $\Pi(e)$ is a strictly concave function of e since $c(e)$ is strictly convex. Given the iñada conditions on $\Pi(e)$ (through $c(e)$), the applicant's program has an interior solution. Thus, given e^*, the solution is uniquely determined by the first order condition

$$\psi(e,e^*) = \gamma p\Delta\overline{V} + (1-p)\gamma[1-E_h(e^*)]\underline{V}_G - (1-p)\gamma[1-E_l(e^*)]\underline{V}'_G - c'(e) = 0.$$

A rational expectation equilibrium is a fixed point of ψ; that is

$$\psi(e^*,e^*) = 0, \tag{17}$$

$$\text{or } \gamma p\Delta\overline{V} + (1-p)\gamma[1-E_h(e^*)]\underline{V}_G - (1-p)\gamma[1-E_l(e^*)]\underline{V}'_G = c'(e^*). \tag{18}$$

Existence is guaranteed by the iñada conditions and the continuity of $\Pi(e)$. To study uniqueness, we use the equilibrium condition (8) and we compute $\partial\psi(e,e)/\partial e = -(1-p)\gamma\left[\underline{V}_G - \underline{V}'_G/(1-\alpha)\right]\partial E_h(e)/\partial e - c''(e)$. From the proof of Proposition 2, $\partial E_h(e)/\partial e > 0$. It thus follows that $\psi(e,e)$ is strictly monotonic (decreasing) if $\alpha \leq 1 - \underline{V}'_G/\underline{V}_G$.

To study when e^* is positive, we first compute the value of $\psi(e^*,e^*)$ at the boundaries. When $\alpha = 0$, $E_h(e^*) = E_l(e^*) = E^* > 0$ and condition (18) is $\psi(e^*,e^*) = \gamma p\Delta\overline{V} + (1-p)\gamma(1-E^*)\Delta\underline{V} - c'(e^*) = 0$, which yields a strictly positive applicant's search effort.

When α tends toward α_1, θ^* tends toward 1, and $E_h(e^*)$ tends toward E_h^t and $E_l(e^*)$ tends toward E_l^t. Condition (17) is then $\psi(e^*,e^*) = \gamma p\Delta\overline{V} + (1-p)\gamma\left[1-E_h^t\right]\underline{V}_G - (1-p)\gamma\left[1-E_l^t\right]\underline{V}'_G - c'(e^*) = 0$. Using the optimal effort levels, this condition yields a strictly positive effort if $\gamma p\Delta\overline{V} + (1-p)\gamma(\Delta\underline{V} - (1-p)\Delta W\left(\underline{V}_G - x_l\underline{V}'_G\right)) > 0$, which is (11).

To show the monotonicity with respect to α, we rewrite (18) as $\gamma p\Delta\overline{V} + (1-p)\gamma(1-E_h^*)\underline{V}_G - (1-p)\gamma(1-E_l^*)\underline{V}'_G = c'(e^*)$. For a given e^* the derivative of the left-hand side with respect to α is negative as $\partial E_h^*/\partial\alpha > 0$ and $\partial E_l^*/\partial\alpha < 0$ (by using the proof of Lemma 3). Thus, $\psi(e^*,e^*)$ is a decreasing function of α.

When condition (18) is not satisfied, the left-hand side of (18) is strictly negative for $\alpha = \alpha_1$. Thus, given continuity and monotonicity, there exists $\widetilde{\alpha} < \alpha_1$ such that the left-hand side is 0. For $\alpha > \widetilde{\alpha}$, the applicant exerts no effort.

Proof of Propositions 4 and 5   This proof is deduced from the proofs of Lemma 3 and Proposition 2. With the SPER, θ_{SR}^* is the solution of $E_h(\theta) = \beta(1-\alpha) + (1-\beta)(1-\alpha)E_l(\theta)$. These functions are represented in Fig. 2. By using Fig. 2, we see that the effect of β is equivalent to the effect of α for $\alpha < \alpha_1$. Therefore, θ_{SR}^* as the solution of (12) is such that $\theta_{SR}^* > \theta^*$, $E_h^{SR} > E_l^{SR}$, $E_h^{SR} > E_h^*$ and $E_l^{SR} < E_l^*$.

The applicant's effort is determined by using Eq. (10) with the appropriate level of effort E_h^{SR}. By using Eq. (10) for E_h^{SR} and E_h^*, we show that $e_{SR}^* < e^*$ for $\alpha = 0$.

Proof of Lemma 4   Due to less strategic reporting from bad applicants, for any $\beta > 0$, a second review increases (resp., reduces) the number of applications submitted with $\tilde{x} = 1$ (resp., $\tilde{x} = x_l$). As $\mu_h(\theta^{SR}) < \mu_h(\theta^*)$, the total number of patents that is granted is smaller in mature fields.

We then study how the composition of the applications changes when β increases once the effort e has been performed. The number of applications with high prior art content ($\tilde{x} = 1$) that are scrutinized is $N_s^h = e\gamma[p + (1-p)\theta]$, while the number of applications with low prior art content ($\tilde{x} = x_l$) is $N_s^l = e\gamma(1-p)(1-\theta) + 1 - e\gamma$. Given that $\partial\theta/\partial\beta > 0$, we obtain $\partial N_s^h/\partial\beta > 0$ and $\partial N_s^l/\partial\beta > 0$.

For abundant prior art, the number of granted patents is $N_g^h = e\gamma[p + (1-p)\theta(1-E_h)]$, and for scarce prior art it is $N_g^l = (1-e\gamma)p + e\gamma(1-p)(1-\theta)(1-E_h) + (1-e\gamma)(1-p)(1-E_h)$.

We compute that $\partial N_g^h/\partial\beta > 0$ and $\partial N_g^l/\partial\beta < 0$. The ratios of good patented innovations over all patented innovations for abundant prior art and scarce prior art are $R^h = (e\gamma p)/N_g^h$ and $R^l = (1-e\gamma)p/N_g^l$. Comparative statics follow from the previous calculations.

Proof of Proposition 6   Let $A = (E_h^{cA}, E_l^{cA})$ be a policy such that $E_l^{cA} < E_h^{cA}$. In such a policy, the equilibrium beliefs of the examiner are $\mu_h^{cA} = 1$ and $\mu_l^{cA} = p(1-\gamma)/(1-\gamma p) = \mu_l < p$.

Now consider a policy $A' = (E_h^{cA'}, E_l^{cA'})$, such that $E_l^{cA'} = E_l^{cA}$ and $E_h^{cA'} = E_h^{cA} - \varepsilon$, with $\varepsilon > 0$ but with $E_h^{cA'} > E_l^{cA'}$. With this new policy, upon receiving $\tilde{x} = 1$, the examiner exerts efforts $E_h^{cA'} > E_l^{cA'}$. Thus, a bad applicant will never choose to report $\tilde{x} = 1$ and $\mu_h^{cA'} = 1$. A bad applicant will report $\tilde{x} = x_l$, so that $\mu_l^{cA'} = p(1-\gamma)/(1-\gamma p) < p$.

The examiner exerts effort $E_h^{cA'}$, pays a cost $C(E_h^{cA'})$, and follows her beliefs. It is obvious that policy A' dominates policy A, since they have the same screening efficiency. But A is more costly than A', as $C(E_h^{cA'}) < C(E_h^{cA})$, even though $C(E_l^{cA'}) = C(E_l^{cA})$.

Proof of Proposition 7   By using the first-order condition (16), we define $F(E^c)$ as

$$F(E^c) \equiv (1-p)x_l\Delta W - \left[1 + (1-p)(1-x_l)\gamma\frac{\Delta V}{c''(e^c)}\frac{E^c}{2}\right]E^c + (1-x_l)\gamma e^c E^c = 0.$$

To study the existence and the number of roots of $F(.)$, we first compute its value at the boundaries. Thus $F(0) \equiv (1-p)x_l\Delta W > 0$; and, if we denote by $\overline{e}^c$ the solution

of (14) when $E^c = 1$, $F(1) < 0$ as $(1-p)\Delta W < 1$ and $\bar{e}^c \leq 1$. Hence, Eq. (16) has at least one root in E^c between 0 and 1. To show uniqueness, we compute the first derivative of $F(E^c)$,

$$\frac{dF(E^c)}{dE^c} = -(1 - \gamma(1-x_l)e^c) - \frac{(1-p)\gamma\Delta\underline{V}}{c''(e^c)}E^c + \gamma E^c\left[1 - x_l + (1-p)\Delta\underline{V}\frac{c'''(e^c)}{2[c''(e^c)]^2}E^c\right]\frac{de^c}{dE^c}.$$

This derivative is negative given iñada conditions on the applicant's effort cost and $c'''(.) \geq 0$. To discuss comparative statics, we calculate $dF(E^{c*})/d\Delta W = (1-p)x_l > 0$, and

$$\frac{dF(E^{c*})}{d\Delta\overline{V}} = (1 - x_l)\gamma E^c\left[(1-p)\Delta\underline{V}\frac{c'''(e^c)}{2[c''(e^c)]^2}E^c + 1\right]\frac{de^c}{d\Delta\overline{V}} > 0.$$

Proof of Proposition 8 In the commitment case, for any given E^c, the applicant's program is

$$\max_e \Pi^c(e) = e\gamma[p\overline{V}_G + (1-p)(1-E^c)\underline{V}_G] + (1-e)\gamma[p\overline{V}'_G + (1-p)(1-E^c)\underline{V}'_G]$$
$$+ (1-\gamma)[p\overline{V}_G + (1-p)(1-E^c)\underline{V}_G] - c(e),$$

and, thus, the optimal effort of the applicant, e^c, is the solution of

$$\gamma p\Delta\overline{V} + (1-E^c)(1-p)\gamma\Delta\underline{V} = c'(e^c). \tag{19}$$

In the non-commitment case, for $\alpha = 0$ the optimal effort of the applicant, e^*, is the solution of

$$\gamma p\Delta\overline{V} + (1-E^*_h)(1-p)\gamma\Delta\underline{V} = c'(e^*). \tag{20}$$

These two conditions are identical; thus, whenever $E^c \lesseqgtr E^*_h$, $e^c \gtreqless e^*$. Hence, we need to compare the PTO's first-order conditions.

In the commitment case, we study whether $F(E) \equiv (1-p)x_l\Delta W - (1 - \gamma(1-x_l)e)E - (1-p)(1-x_l)\gamma E^2\Delta\underline{V}/2c''(e)$ is strictly negative or strictly positive in equilibrium. In equilibrium $E^*_h = E^*_l = E^*$; equivalently, $e^* = (\theta^*(1-x_l) - px_l(1-\theta^*))/(1-x_l)\theta^*\gamma[p + \theta^*(1-p)]$ by using expressions of E^*_h and E^*_l (from proof of Lemma 3). Given that $e^* > 0$, it implies that $\theta^* > \underline{\theta} \equiv px_l/(1 - x_l + px_l)$. Using e^* and E^*_l we obtain

$$F(E^*) = (1-p)^2(1-\theta^*)\frac{\theta^*(1-x_l) - px_l(1-\theta^*)}{(1-x)\theta^*\gamma[p + \theta^*(1-p)]} - (1-p)(1-x_l)\gamma\frac{\Delta\underline{V}}{c''(e^*)}\frac{(E^*_l)^2}{2}. \tag{21}$$

For any interior solution, $\theta^* > \underline{\theta}$, and the first term of (21) is strictly positive. We now study the sign of $F(E^*)$ when $\Delta\underline{V}$ varies on $\mathbb{R}_+$. For any $\Delta\overline{V} > 0$,

Table 3 Summary of comparative statics

Parameter y	$\frac{\partial \alpha_1}{\partial y}$	$\frac{\partial E_h^l}{\partial y}$	$\frac{\partial E_l^l}{\partial y}$	$\frac{\partial \theta^*}{\partial y}$	$\frac{\partial E_h^*}{\partial y}$	$\frac{\partial E_l^*}{\partial y}$	$\frac{\partial E_h^{SR}}{\partial y}$	$\frac{\partial E_l^{SR}}{\partial y}$	$\frac{\partial E_c}{\partial y}$
p	−	−	−	?	−	−	−	−	
γ	.	.	.	+	+	+	+	+	
α	.	.	.	+	+	−	+	−	
ΔW	+	+	+	?	+	+	+	+	+
x_l	−	.	+	+	+	+	+	+	
e	.	.	.	+	+	+	+	+	
β	.	.	.	.	.	.	+	+	

$$F(E^*) \underset{\Delta \underline{V} \to 0^+}{\approx} (1-p)^2 (1-\theta^*) \frac{\theta^*(1-x_l) - p x_l(1-\theta^*)}{(1-x_l)\theta^*\gamma[p + \theta^*(1-p)]} > 0,$$

since $c''(e^*)$ and E_l^* are finite (e^* is defined by $c'(e^*) = \gamma p \Delta \overline{V}$). Hence, taking into account the strict monotonicity of $F(.)$, $F(E^*) > 0$ implies that $E^* < E^c$ when $\Delta \underline{V} \to 0^+$. Using (19) and (20), $e^c < e^*$. This proves the existence of a threshold $\underline{\Delta}$.

When $\Delta \underline{V} \to \infty$, we check that $F(E^*) < 0$. Using identical arguments as above, $E^* > E^c$ when $\Delta \underline{V} \to \infty$. Using (19) and (20), $e^c > e^*$. The existence of a threshold $\overline{\Delta}$ follows in Table 3.

References

Aghion, P., & Tirole, J. (1997). Formal and real authority in organizations. *Journal of Political Economy*, *105*(1), 1–29.

Alcácer, J., & Gittelman, M. (2006). Patent citations as a measure of knowledge flows: The influence of examiner citations. *The Review of Economic Statistics*, *88*(4), 774–779.

Alcácer, J., Gittelman, M., & Sampat, B. (2009). Applicant and examiner citations in U.S. patents: An overview and analysis. *Research Policy*, *38*(2), 415–427.

Allison, J., & Hunter, S. (2006). On the feasibility of improving patent quality one technology at a time: The case of business methods. *Berkeley Technology Law Journal*, *21*, 729–794.

Allison, J., & Lemley, M. (1998). Empirical evidence on the validity of litigated patents. *American Intellectual Property Law Association (AIPLA) Quarterly Journal*, *26*, 185.

Atal, V., & Bar, T. (2010). Prior art: To search or not to search. *International Journal of Industrial Organization*, *28*, 507–521.

Atal, V., & Bar, T. (2014). Patent quality and a two-tiered patent system. *The Journal of Industrial Economics*, *62*(3), 503–540.

Chang, H. (1995). Patent scope, antitrust policy and cumulative innovation. *RAND Journal of Economics*, *26*, 34–57.

Chiou, J.-Y. (2008). The patent quality control process: Can we afford an (rationally) ignorant patent offices? CLEA Meetings Paper.

Cockburn, I, Kortun, S., & Stern, S. (2003). Are all patent examiners equal? Examiners, patent characteristics and litigation outcomes. In W. Cohen (Ed.), *Patents in the knowledge-based economy*. Washington: National Academy of Sciences.

Crampes, C., & Langinier, C. (2002). Litigation and settlement in patent infringement cases. *RAND Journal of Economics*, *33*(2), 258–274.

Farrell, J., & Merges, R. (2004). Incentives to challenge and defend patents: Why litigation won't reliably fixx patent office errors and why administrative patent review might help? *Berkeley Technology Law Journal*, *19*(1), 2–28.

Federal Trade Commission. (2003). To promote innovate: The proper balance of competition and patent law ans policy. Report.

Friebel, G., Koch, A., Prady, D., & Seabright, P. (2006). Objectives and incentives at the European patent office. IDEI Report.

Graham, S., Hall, B., Harhoff, D., & Mowery, D. (2002). *Post-issue patent "quality control:" A comparative study of U.S. patent re-examinations and European patent oppositions.* Cambridge: National Bureau of Economic Research.

Jaffe, A., & Lerner, J. (2004). *Innovation and its discontents: How our broken patent system is endangering innovation and progress, and what to do about it.* Princeton: Princeton University Press.

Kesan, J. (2002). Carrots and sticks to create a better patent system. *Berkeley Technology Law Journal*, *17*, 145–179.

Khalil, F. (1997). Auditing without commitment. *RAND Journal of Economics*, *28*, 625–640.

Lampe, R. (2012). Strategic citation. *Review of Economics and Statistics*, *94*(1), 320–333.

Langinier, C., & Marcoul, P. (2014). Subject performance of patent examiners and implicit contracts. Working paper.

Lanjouw, J., & Schankerman, M. (2001). Characteristics of patent litigation: A window on competition. *RAND Journal of Economics*, *32*, 129–151.

Lemley, M., Lichtman, D., & Sampat, B. (2005). What to do about bad patents? *Regulation*, *2005–2006*, 10–13.

Lerner, J. (2002). Where does state street lead? a..rst look at..nance patents, 1971 to 2000. *The Journal of Finance*, *LVIII*, 901–930.

Levitt, S., & Snyder, C. (1997). Is no news bad news? Information transmission and the role of "early warning" in the principal–agent model. *RAND Journal of Economics*, *28*, 641–661.

Merges, R. (1999). As many as six impossible patents before breakfast: Property rights for business concepts and patent system reform. *Berkeley Technology Law Journal*, *14*, 577–615.

Mitusch, K. (2006). Non-commitment in performance evaluation and the problem of information distortions. *Journal of Economics Behavior and Organization*, *60*, 507–525.

O'Donoghue, T. (1998). A patentability requirement for sequential innovation. *RAND Journal of Economics*, *29*, 654–679.

Picard, P., van Potterie, B., & Pottelsberghe de la. (2013). Patent office governance and patent examination quality. *Journal of Public Economics*, *104*, 14–15.

Quillen, C., & Webster, O. (2001). Continuing patent applications and performance of the U.S. patent office. *Federal Circuit Bar Journal*, *11*(1), 1–21.

Quillen, C., & Webster, O. (2006). Continuing patent applications and performance of the U.S. patent and trademark office-updated. *The Federal Circuit Bar Journal*, *15*(4), 635–677.

Sampat, B. (2010). When do applicants search for prior art. *Journal of Law and Economics*, *53*(2), 399–416.

Sampat, B., & Lemley, M. (2010). Examining patent examination. *Stanford Technology Law Review*, *2010*, 2.

Scotchmer, S. (1996). Protecting early innovators: Should second generation product be patentable? *RAND Journal of Economics*, *27*, 322–331.

Strausz, R. (2005). Timing of verification procedures: Monitoring versus auditing. *Journal of Economic Behavior and Organization*, *59*, 89–107.

Trajtenberg, M., Jaffe, A., & Fogarty, M. (2000). Knowledge spillovers and patent citations: Evidence from a survey of inventors. *American Economic Review, Papers and Proceedings*, *90*, 215–218.