Inventors and Impostors: An Economic Analysis of Patent Examination*

Florian Schuett[†]

November 2008

Abstract

We present a model in which firms differing in creativity decide whether to invest in genuine research or to submit "bogus" patent applications (claiming that they have invented something which is not truly novel). The government delegates the verification of novelty to an agency which must exert costly effort in order to obtain a signal of patentability. Firms self-select according to their creativity, with high-creativity types producing true innovations and low-creativity types submitting bogus applications, or staying idle. The thresholds depend on the expected examination effort and on the application fee. We show that, at the full-commitment optimum, all bogus applications are deterred. When the agency lacks commitment power and its effort is unobservable, the outcome hinges on whether the patentability signal is hard or soft information. With hard information, the government rewards the agency for rejecting patent applications and can attain an allocation that is arbitrarily close to the optimum. With soft information, however, transfers are constrained by the need to ensure truthfulness, creating a tradeoff between innovation and deterrence.

Keywords: innovation, patent office, limited commitment, incentives for bureaucrats

JEL classification numbers: O31, O38, D73, D82, L50

* I thank Jing-Yuan Chiou, Pascal Courty, Vincenzo Denicolò, Guido Friebel, Elisabetta Iossa, Paolo Pin, François Salanié, Mark Schankerman, Paul Seabright, Jean Tirole, participants at the EVPAT summer school in Bologna, the "Knowledge for Growth" conference in Toulouse, the congress of the European Economic Association in Milan, the EARIE conference in Toulouse, the EPIP conference in Bern, as well as seminar participants at Toulouse School of Economics and the European University Institute for helpful comments and suggestions. All errors are mine.

† Toulouse School of Economics (GREMAQ) and European University Institute. Postal address: Villa La Fonte, Via delle Fontanelle 10, 50014 San Domenico, Italy. Email: florian.schuett@eui.eu

1 Introduction

Because of the inefficiencies of monopoly, patents should only be granted for true inventions. Granting patents for non-inventions causes deadweight loss and litigation without providing any offsetting benefit to society. If anticipated, it can also divert resources from research to rent-seeking. The patent office plays the role of a watchdog making sure that only novel and non-obvious inventions obtain patent protection. Yet, as infringement lawsuits filed by holders of dubious patents against prominent firms such as eBay and RIM have brought to public awareness, the patent office does not reliably weed out bad patents.[1] The cost of the bad patents that slip through the net has been estimated at an annual $25.5 billion for the US economy (Ford *et al.*, 2007).[2]

To many observers, the failure of the patent office to rigorously screen patent applications is a source of concern.[3] Lemley (2001), however, argues that the patent office may rationally choose to spend limited resources on examining a given application because only a tiny fraction of patents ever turn out to be commercially significant. The cost of screening out more bad applicants might well exceed the benefit. In this paper, we show that this argument misses two important points. First, it considers only the ex post benefits of examination. But when firms choose which activity to pursue depending on how rigorous they expect patent examination to be, ex ante benefits may exceed ex post benefits. Second, it takes the application fee as given. But under an optimal patent policy, the application fee and examination effort should be jointly determined.

The objectives of the paper are to characterize an optimal patent policy, and to offer an explanation of why it may differ from observed policy. We develop a model in which firms differing in creativity self-select to become either inventors doing research or impostors submitting bogus applications. In this setup, greater examination effort increases the share of firms doing research. We show that the optimal policy involves full deterrence of bogus applications. Effort is chosen to balance the benefits of research with the costs of patent examination, while the application fee is used to achieve deterrence.

To address the question of why we don't observe this outcome in the real world, we start from the idea that patent examination resembles an inspection game and as such is

[1] For example, RIM (Research In Motion), the maker of BlackBerry mobile devices, was sued by patent-holding company NTP, and settled out of court for a reported $612.5 million, even though on re-examination the US Patent and Trademark Office (USPTO) revoked all of the patents NTP had asserted against RIM. See Time Magazine, "Patently Absurd", April 2, 2006, available online at http://www.time.com/time/magazine/article/0,9171,1179349,00.html.

[2] Of this sum, $4.5 billion is attributable to litigation costs while the remainder corresponds to the disincentive to future innovators that patents create. While methodologically controversial, Ford *et al.*'s (2007) calculations indicate that the costs of bad patents are likely to be significant.

[3] See, e.g., Jaffe and Lerner (2004), Farrell and Shapiro (forthcoming) and Bessen and Meurer (2008).

plagued by commitment problems. Full deterrence makes examination ex post inefficient and eliminates the incentive for the patent office to perform rigorous screening. It turns out, however, that lack of commitment does not hurt efficiency much when the signal produced by patent examination is hard information. Explaining the observed laxity of examination instead requires a combination of commitment problems and soft information. We find that, with soft information, there is a tradeoff between deterrence and innovation. In addition, the patent office's intrinsic motivation, while irrelevant with hard information, plays a crucial role when the signal is soft information.

In the model presented in section 2, the government delegates patent examination to an agency motivated by both extrinsic rewards (i.e., monetary transfers) and intrinsic rewards (defined as a concern for social welfare). The agency must expend effort to obtain a signal about patentability. If a claimed invention is not truly new, the agency can come up with prior art demonstrating the lack of novelty. Firms differ in their ability to produce valuable inventions (their creativity) and choose whether to do genuine research or to file bogus applications on existing technologies, hoping to escape detection by the patent office. While genuine research creates value for society, granting monopoly power to impostors causes social losses.[4] The private profitability of the two activities depends on the patent office's examination effort. More rigorous examination makes it less likely for impostors to get away, and therefore increases the attractiveness of true research. This setup leads to self-selection of firms: under a single-crossing condition, high-creativity firms do genuine research, while low-creativity firms submit bogus applications or stay idle. Our formulation acknowledges that the patent office may have a role in encouraging R&D, as stressed by Jaffe and Lerner (2004).

The government chooses an application fee and an incentive scheme for the patent office. While it could deter all bogus applications solely through a sufficiently large application fee and thereby avoid the cost of patent examination, such a policy leads to a suboptimal level of innovation. We show in section 3 that under the optimal full-commitment policy, effort is chosen so as to equalize the marginal gains from innovation with the marginal cost of patent examination. The application fee is set at the level that deters all bogus applications. Thus, at the optimum, no invalid patents are issued. In the absence of commitment, this feature of the optimal policy creates problems. If the patent office expects all applications to be valid, it has no incentive to exert effort. But if the patent office's examination effort is low, some firms are better off submitting bogus applications rather than actually doing research.

[4] In practice, the possibility of challenging a patent in court mitigates this problem, but does not eliminate it if the court decision is uncertain. There may also be too little challenging of questionable patents because of the public good nature of these challenges (Chiou, 2006; Farrell and Shapiro, forthcoming). In any case, the costs of patent litigation are substantial in their own right.

We then study the outcome of the examination game between firms and the patent office when the office lacks commitment and its effort is unobservable. Section 4 establishes our main results. We distinguish two cases according to whether the patent office's signal is hard or soft information. When the signal is hard information, we show that the government can design an incentive scheme that allows it to come arbitrarily close to the optimal full-commitment solution. The scheme rewards the patent office for coming up with prior art that is grounds for rejecting an application. The application fee is adjusted to achieve the optimal amount of innovation. The agency's level of intrinsic motivation does not affect the outcome.

With soft information, the sunk-cost nature of R&D investment and the deadweight loss from monopoly pricing create a second commitment problem – namely over the grant decision. The agency will be tempted to reject even valid applications in order to avoid deadweight loss. Incentives must be designed to make sure that rejections only occur when the agency has actually found invalidating prior art, which requires that it be paid for granting patents. Since the incentive scheme must be used to ensure truthfulness on the part of the agency, the government's only remaining effective instrument is the application fee. As a result, the full-commitment outcome can no longer be attained. In choosing an application fee, the government trades off the benefits from innovation against the costs of invalid patents. A lower application fee leads to more bogus applications but at the same time incites the patent office to screen more rigorously, which, in turn, leads to more research. Moreover, inducing both truthful revelation and effort provision requires intrinsic motivation on the part of the agency.

We then discuss whether hard or soft information is a better description of reality and compare the predictions of the model with empirical observations. We argue that the complexity of patent applications and the inherent vagueness of the non-obviousness standard confer considerable discretion on patent examiners. This suggests that it may be more appropriate to consider prior art as soft information. The soft-information model also produces results which are more in line with what we observe in practice. It can account for the apparent laxity of patent examination reflected by the high incidence of bad patents. The incentive scheme that ensures truthfulness in the soft-information case is also roughly consistent with compensation practice at the USPTO, where patent examiners are paid for the number of applications treated. Combined with rules that make rejections more time-consuming than grants, this amounts to a bonus for accepting applications (Merges, 1999). Finally, unlike in the hard-information case, the efficiency of patent examination depends on the agency's concern for social welfare. Such intrinsic motivation has been identified as an important characteristic of many bureaucracies (Perry and Wise, 1990), and of examiners at the European Patent Office (EPO) in particular (Friebel *et al.*, 2006).

In section 5, we check the robustness of the results to alternative assumptions. We first look at the implications of introducing a shadow cost of public funds. While both the full-deterrence result in the full-commitment case and the attainability of optimum in the hard-information case may seem to rely on the absence of such a cost, we show that under plausible conditions, both results are robust. We also investigate how results change when the ex post welfare effects of patents are positive. The planner may no longer reward the agency for grants when information is soft, but truthful revelation continues to put an upper bound on the power of incentives.

We then develop a more structural model of the innovation process and identify the conditions under which the assumptions on profit and welfare functions are likely to hold. Building on the structural model, we examine two extensions that enlarge the planner's set of instruments. The first allows the planner to use probabilistic patent grants, while the second introduces inadvertent re-invention and lets the planner differentiate fees according to the signal. Finally, section 6 summarizes the main findings, discusses limitations, and comments on policy implications.

Related literature

The paper contributes to the literature on optimal patent policy (see, e.g., Cornelli and Schankerman, 1999; Scotchmer, 1999; Hopenhayn and Mitchell, 2001; Hopenhayn *et al.*, 2006). In taking a mechanism-design approach, this literature implicitly assumes that for policy implementation, the mechanism designer can rely on a benevolent agency with commitment power. Moreover, while the literature assumes that the inventor has private information on the value of his innovation, it does not examine the novelty dimension. However, like any regulator, the patent office may be tempted to engage in opportunistic behavior.[5] If the government needs to give discretion over patenting decisions to an agency having to exert effort to learn about the novelty of applications, then opportunism with respect to both examination effort and the grant decision arises naturally.

The paper is also related to the auditing literature, which has addressed similar commitment problems and documented that the nature of the signal produced by the audit is crucial. For the hard-information case, Melumad and Mookherjee (1989) show that delegation of auditing decisions to a salary-maximizing manager can achieve the full-commitment outcome, a result we extend to the case of patent examination. Iossa and Legros (2004) study auditing with soft information and show that a necessary condition for the auditor to exert any effort is that he be given a stake in the audited project. Similarly, we show in the soft-information

[5] See Armstrong and Sappington (2007) for an overview of the issues that arise under limited regulatory commitment.

case that positive effort will only occur if the examiner is intrinsically motivated – that is, if he has a "stake" in the resulting level of social welfare.

A number of recent papers investigate patent examination. Prady (2008) and Langinier and Marcoul (2003) consider signaling models. Prady (2008) shows that leniency on the part of the patent office may be necessary to induce true inventors to fully describe their invention, which is assumed to increase welfare. Thus, she offers an alternative argument to ours as to why patent examination may be lax. Langinier and Marcoul (2003) study inventors' incentives to search for and disclose relevant prior art to the patent office. They find that, when the patent office cannot commit to a level of screening, there exists no equilibrium where applicants having obtained a positive signal separate from applicants with a negative signal in terms of the amount of prior art they submit. This is because, as in our model, the patent office has no incentive to search if it can identify valid applications beforehand.

Caillaud and Duchêne (2005) present a model in which valid inventions stem from successful R&D projects and invalid ones from failed projects. They focus on the "overload problem" facing the patent office: when flooded with large numbers of applications, the average quality of examination declines, leading to a vicious circle by encouraging even more invalid applications. Again, there cannot be a separating equilibrium, i.e., one where only valid applicants file for a patent. Both Langinier and Marcoul (2003) and Caillaud and Duchêne (2005) assume that the patent office lacks commitment power but is fully benevolent. Thus, they do not take into account the result obtained by Melumad and Mookherjee (1989), according to which delegation to an agency that responds to monetary incentives can overcome the commitment problem. Moreover, their analysis relies on the patent office's signal being hard information, while we consider both the hard and soft-information cases and show that they lead to very different conclusions.

Chiou (2008) looks at how the patent office's examination effort interacts with the incentives of private parties to bring court challenges. He shows that the two types of enforcement may be complementary, and that weak patents are more likely to be settled out of court than strong patents. His results, as ours, call into question Lemley's (2001) rational ignorance argument. In his model, however, the patent office has full commitment power. Finally, in none of these models do inventors choose whether to do genuine research or not, so unlike in our model, the proportion of potential bad applicants is exogenous.

2 The model

Consider the following setup. There are three types of players: a benevolent planner (the government or Congress), an agency (the patent office), and a mass 1 of firms. Firms are characterized by a creativity parameter θ which is their private knowledge and distributed according to cdf $G(\cdot)$ on $[0, \infty)$.

Assumption 1. *The distribution of θ satisfies the monotone hazard rate property:*

$$\frac{\mathrm{d}}{\mathrm{d}\theta}\left(\frac{g(\theta)}{1-G(\theta)}\right) \geq 0.$$

Firms are endowed with one indivisible unit of time which they can devote either to R&D or to filing a bogus patent application claiming something that is either obvious or not novel. Alternatively, firms can stay idle. The idea is that there are existing technologies or obvious combinations of existing technologies that (a) firms can claim to have invented and which are not easily distinguishable from true inventions, and that (b), if awarded a patent, allow the patent holder to extract rents from users; a necessary condition is that such bad patents are enforced by the courts with positive probability. Denote a firm's decision by $d(\theta) \in \{R, B, I\}$. If it does R&D ($d(\theta) = R$), its payoff in case it is awarded a patent is $\pi_R(\theta)$.[6] If it submits a bogus application ($d(\theta) = B$) and obtains a patent, its payoff is $\pi_B(\theta)$ (which can be thought of as the expected profit taking into account that the patent may be invalidated by the courts later on). We assume for simplicity that firms make zero profit if they fail to obtain a patent. Their payoff when staying idle ($d(\theta) = I$) is also zero.

The planner, whose objective is to maximize social welfare, delegates patent examination to the agency. The agency does not observe an applicant's activity (R or B) but does receive a signal σ. The distribution of the signal depends on the applicant's activity and the agency's unobservable examination effort. If the application is for a genuine invention (R), there is no signal ($\sigma = \varnothing$). For bogus applications (B) the patent office obtains a signal indicating that the application is bogus ($\sigma = B$) with probability e, and no signal with probability $1-e$, where $e \in [0, 1]$ is the effort that the agency puts into patent examination. An important question is whether $\sigma = B$ is hard information, in the sense of being verifiable by third parties; we will be more specific on this in the following sections. The cost of effort is $\gamma(e)$ per application that is examined (γ increasing and convex with $\gamma(0) = \gamma'(0) = 0$ and $\gamma'(1) = \infty$).

The agency has utility

$$U = \alpha \cdot [\text{welfare}] + (1-\alpha) \cdot [\text{transfers}] - \gamma(e) \cdot [\text{number of applications}]$$

and is protected by limited liability (i.e., transfers must be nonnegative). The parameter $\alpha \in (0, 1)$ measures how much the agency cares about society's well-being relative to transfers.[7] Importantly, we assume that *both* matter to the agency. We will refer to α as the

[6] This can be seen as a reduced-form profit function resulting from a firm's investment choice; see footnote 11 below and section 5.3.

[7] The welfare variable entering the agency's utility function excludes the agency's own welfare.

agency's level of intrinsic motivation. While the economic literature has only recently begun to acknowledge the importance of intrinsic motivation for understanding bureaucracies,[8] in the public administration literature the concept of "public-service motivation" has a long tradition, and its relevance is empirically established (Perry and Wise, 1990).

Genuine innovations generate social welfare (profits plus consumer surplus) $W(\theta)$ if a patent is issued, with $W' \geq 0$. In the absence of patent protection, social welfare is $W(\theta)+D$. Thus, D is the difference between the social surplus the invention generates with and without patent protection. Such a difference may arise for several reasons which we discuss in sections 4.3 and, more formally, 5.3 below. A patent on a bogus application causes a social loss of $L > 0$.

Assumption 2. *The ex post welfare effects of patents satisfy*

$$L > D \geq 0.$$

The first inequality in Assumption 2 ($L > D$) says that bad patents cause greater social losses than valid patents. The second inequality ($D \geq 0$) says that patents never increase welfare from an ex post point of view; this is consistent with the standard view of patents as providing incentives for invention at the cost of reducing their use. For simplicity, D and L are also assumed to be independent of θ.[9]

The timing of the game is as follows (see figure 1). At the beginning of the game, the planner sets an application fee F and chooses an incentive scheme for the agency.[10] Then, firms choose their activity $d(\theta)$. The agency decides how much examination effort e to provide. Finally, signals are drawn, acceptance and rejection decisions are made, and payoffs are realized. The important assumption here is that the agency cannot commit to a level of examination effort e before firms choose their activity.

Firm behavior

Given an application fee F and an anticipated examination effort e, each type of firm chooses $d(\theta)$ to maximize its expected payoff. Genuine research yields a payoff of $\pi_R(\theta) - F$, while a bogus applicant can expect net profit $(1-e)\pi_B(\theta) - F$.[11] Thus, a firm prefers R&D

[8] See, e.g., Prendergast (2007).

[9] Section 5.3 develops a more structural model of the innovation process. The model accounts in a simple way for secrecy, sequential innovation, and court challenges. It allows us to derive conditions under which welfare functions satisfy the assumed properties, in particular Assumption 2.

[10] This is a restriction on the set of instruments that the planner has at her disposal. In particular, we don't allow grant probabilities other than 0 and 1, i.e., the grant decision is deterministic, and we impose a uniform application fee instead of conditioning fees on the outcome of the examination. We consider these extensions in section 5.4.

[11] If $\pi_R(\theta)$ is interpreted as a reduced-form profit function resulting from the firm's investment choice, one may wonder whether examination effort and the application fee influence the optimal R&D investment, which

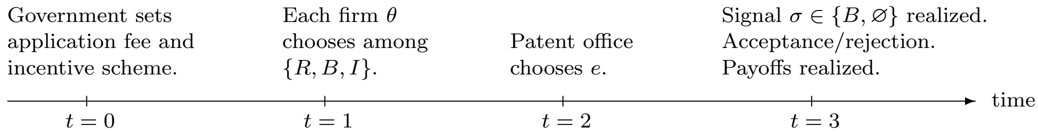

Figure 1: Timing of the game

to imposture if and only if

$$\pi_R(\theta) \geq (1 - e)\pi_B(\theta).$$

Assumption 3. *Profit functions satisfy*

(i) $\pi'_R > \pi'_B > 0$,

(ii) $\pi_R(0) < \pi_B(0)$ *and* $\pi_B(0) \geq 0$,

(iii) $\lim_{\theta\to\infty} \pi_R(\theta) = \infty$ *or* $\lim_{\theta\to\infty}[\pi_R(\theta) - \pi_B(\theta)] > 0$,

(iv) $\pi''_R \leq 0$ *and* $\pi''_B \geq \pi''_R$.

Profits from both activities increase with θ, perhaps because identifying valuable bogus applications requires some of the same qualities as identifying valuable research projects. Profits from research are more sensitive to creativity than those from bogus patents, though. For firms at the lower end of the creativity distribution ($\theta = 0$), obtaining a patent on a bogus application is more profitable than producing a true invention, while towards the upper end of the distribution, it is the opposite. Finally, first derivatives of the profit functions satisfy monotonicity conditions.[12] This "single-crossing" assumption is sufficient for the existence of a unique threshold $\hat{\theta}$ such that, in the absence of application fees, $d(\theta) = B$ for all $\theta < \hat{\theta}$ and $d(\theta) = R$ for all $\theta \geq \hat{\theta}$. The threshold depends on the (expected) effort, i.e., $\hat{\theta} = h(e)$, where h is the implicit function defined by

$$\pi_R(\hat{\theta}) = (1 - e)\pi_B(\hat{\theta}). \tag{1}$$

Moreover, assuming $F \leq \pi_R(\hat{\theta})$, there is a second threshold $\underline{\theta} = \ell(e, F)$ defined by

$$(1 - e)\pi_B(\underline{\theta}) = F, \tag{2}$$

would make the above analysis invalid. However, given the model setup, the level of investment and thus π_R are independent of e and F. To see this, assume (following Cornelli and Schankerman (1999)) that the firm's profit (gross of application fees) is given by $\rho(z, \theta) - \psi(z)$, where z is its R&D investment and $\psi(z)$ the associated cost. Assuming $\rho_z > 0 \geq \rho_{zz}$ (subscripts denote partial derivatives), as well as $\psi' > 0$, $\psi'' > 0$, the optimal amount of R&D effort, $z^*(\theta)$, is determined by $\rho_z(z, \theta) = \psi'(z)$. Clearly, z^* is independent of e and F, and $\pi_R(\theta) = \rho(z^*(\theta), \theta) - \psi(z^*(\theta))$.

[12] Conditions for these assumptions to hold are discussed in section 5.3.

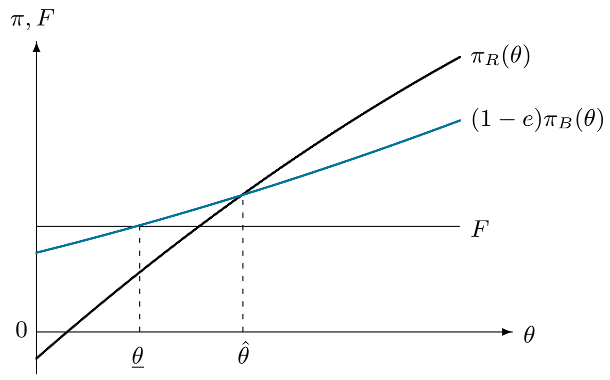

Figure 2: Self-selection of firms according to θ

such that firms with creativity higher than $\hat{\theta}$ do research, firms with creativity between $\hat{\theta}$ and $\underline{\theta}$ submit bogus applications, and firms with creativity lower than $\underline{\theta}$ remain idle. Thus, a patent policy (F, e) leads to self-selection of firms between genuine R&D, imposture, and inactivity, as illustrated in figure 2.

3 Optimal patent policy with full commitment

As a benchmark, we derive the patent policy that the planner would choose ex ante; we will refer to this case as the full-commitment outcome. The optimal combination of e and F maximizes

$$\int_{\hat{\theta}}^{\infty} W(\theta)\mathrm{d}G(\theta) - (1-e)L[G(\hat{\theta}) - G(\underline{\theta})] - \gamma(e)[1 - G(\underline{\theta})] \tag{3}$$

subject to (1), (2) and $\underline{\theta} \leq \hat{\theta}$. The first term corresponds to the social value created by research (undertaken by firms whose creativity exceeds $\hat{\theta}$), the second term captures the expected social losses from issued bad patents, and the third term represents the cost of examination. The constraint $\underline{\theta} \leq \hat{\theta}$ reflects the fact that setting e and F such that $\hat{\theta}$ is strictly below $\underline{\theta}$ can never be optimal. Holding F constant, one could reduce e (and save the associated costs) without changing the set of firms who obtain patents. The following proposition characterizes the optimal patent policy.

Proposition 1. *Suppose Assumptions 1 and 3 hold. The optimal full-commitment policy* (e^o, F^o) *is such that* $\underline{\theta} = \hat{\theta}$. *Examination effort* e^o *satisfies the following equation:*

$$-h'(e^o)W(\hat{\theta})g(\hat{\theta}) = \gamma'(e^o)[1 - G(\hat{\theta})] - h'(e^o)\gamma(e^o)g(\hat{\theta}). \tag{4}$$

The application fee is given by $F^o = \pi_R(h(e^o))$.

Proof: We show first that the constraint $\underline{\theta} \leq \hat{\theta}$ must be binding. Let μ be the multiplier associated with the constraint. Differentiating (3) with respect to F, we have

$$\frac{\partial \ell}{\partial F}\left[g(\underline{\theta})[(1-e)L+\gamma(e)]-\mu\right]=0. \tag{5}$$

Since $\partial \ell/\partial F > 0$, $\mu > 0$, so indeed $\underline{\theta} = \hat{\theta}$. This requires $F = \pi_R(h(e))$. We obtain (4) by differentiating (3) with respect to e, substituting for μ from (5) and using the fact that $\underline{\theta} = \hat{\theta}$. It remains to be shown that the second-order condition holds at e^o, which requires

$$-h''Wg-(h')^2[W'g+Wg']-\gamma''(1-G)+2h'\gamma'g+\gamma\left[h''g+h'g'\right]<0.$$

At e^o, this can be rewritten using the fact that, by (4), $\gamma = \gamma'(1-G)/(h'g) + W$:

$$-(h')^2W'g+(1-G)\left[\frac{h''}{h'}\gamma'-\gamma''\right]+h'\gamma'\frac{2g^2+(1-G)g'}{g}<0.$$

The fraction is positive thanks to Assumption 1. Moreover,

$$h'(e) = -\frac{\pi_B}{\pi_R'-(1-e)\pi_B'} \;\leq\; 0$$

and

$$\begin{aligned} h''(e) &= \frac{\pi_B\left[h'[\pi_R''-(1-e)\pi_B'']+\pi_B'\right]-h'\pi_B'[\pi_R'-(1-e)\pi_B']}{(\pi_R'-(1-e)\pi_B')^2} \\ &= \frac{\pi_B\left[2\pi_B'[\pi_R'-(1-e)\pi_B']-\pi_B[\pi_R''-(1-e)\pi_B'']\right]}{\left(\pi_R'-(1-e)\pi_B'\right)^3} \;>\; 0 \end{aligned}$$

where the inequalities follow from Assumption 3. ■

Higher examination effort increases the attractiveness of genuine research relative to imposture. That is, e determines the incentives to do R&D. The planner chooses e^o to equalize marginal social gains from more innovation (the left-hand side of (4)) with the marginal cost of examination (the right-hand side of (4)). Meanwhile, F^o is set so as to deter all firms with $\theta < \hat{\theta} = h(e^o)$ from applying. At the optimum, there are no bogus applications, and no bad patent is issued. Intuitively, as long as $F < \pi_R(\hat{\theta})$, raising the application fee does not represent a disincentive to innovation in this model: only those types of firm who would anyway find it optimal to submit bogus applications are discouraged from applying for patents. Thus, there is no loss in raising the fee up to the level where imposture is completely deterred.

The fact that the optimal policy is characterized by full deterrence clearly makes it ex post inefficient: the patent office spends resources on patent examination even though all applicants have true inventions. This inefficiency causes problems when examination effort is chosen *after* firms have decided on their activities. Even though the agency cares about welfare (and even if it cared about welfare as much as the planner), it does not take into account the effects of its examination effort on incentives to undertake research; it is only concerned with avoiding issuing invalid patents. The agency will be tempted to cut back on examination effort and screen applications less rigorously than ex ante efficiency would require.

Results from the auditing literature, however, suggest that delegation to an auditor who responds to monetary incentives can solve the principal's commitment problem (Melumad and Mookherjee, 1989; Strausz, 1997). In the following section, we investigate whether delegating patent examination to an agency that cares about both welfare and monetary transfers allows the government to achieve the full-commitment outcome.

4 Incentives for the patent office

This section studies the examination game described in Section 2. The agency lacks commitment power – it chooses effort after firms choose their activity – and its effort is unobservable. We focus on the simple case where there is one patent examiner for every applicant. While not without loss of generality, this assumption has the merit of conveying most of the intuition in a straightforward manner.[13] When we talk about the agency in what follows, it may sometimes be more appropriate to think of a representative examiner; with a slight abuse of language we will use the terms agency and examiner interchangeably.

We start by analyzing the case where the signal $\sigma = B$ is hard information in section 4.1. In section 4.2 we turn to the case where the signal is soft.

4.1 Hard information

We solve the game backwards, starting with the agency's effort choice. Suppose the agency believes that a proportion p of all applicants has patentable inventions.[14] By screening out a bad application, it avoids a social loss of L. The incentive scheme chosen by the planner can be conditioned on two events: either the agency comes up with defeating prior art ($\sigma = B$), in which case it receives transfer t_B and the application is rejected, or it does not ($\sigma = \varnothing$), in which case it receives $t_\varnothing$ and the patent is granted. The agency chooses e to maximize

$$p[(1-\alpha)t_\varnothing - \alpha D] + (1-p)[e(1-\alpha)t_B + (1-e)[(1-\alpha)t_\varnothing - \alpha L] - \gamma(e).$$

[13] The assumption reduces the set of incentive schemes available to the planner. In section 5.1, we study the more general case where each examiner treats many patent applications.

[14] Alternatively, one could assume that the agency observes the number of applications when deciding on e. This would not change the analysis.

With probability p, it encounters a valid applicant, so that it cannot find any grounds for rejection. The transfer it receives is $t_\varnothing$. Since the invention has already been made, the welfare effect of granting a patent is $-D$. With probability $1-p$, the application is bogus, for which the agency finds evidence with probability e. It is paid t_B and the patent is rejected, so no welfare loss is incurred. With probability $1-e$, the agency finds no evidence and the patent is granted, with associated welfare effect of $-L$. Differentiating with respect to e leads to the first-order condition

$$(1-p)\Big[\alpha\, L + (1-\alpha)(t_B - t_\varnothing)\Big] = \gamma'(e). \tag{6}$$

It is obvious from (6) that a strictly positive level of examination effort is only sustainable if the agency expects there to be some bogus applications. The examination game is formally equivalent to an "inspection game". In equilibrium, the applicants' and the agency's choices must be best responses to each other's strategies.

Proposition 2. *There exists a unique equilibrium* (p^*, e^*) *of the examination game such that*

$$p^* = \frac{1-G(h(e^*))}{1-G(\ell(e^*,F))} \tag{7}$$

$$e^* = (\gamma')^{-1}\left((1-p^*)\big[\alpha\, L + (1-\alpha)(t_B - t_\varnothing)\big]\right). \tag{8}$$

For any $F < \pi_R(h(0))$*, the equilibrium is characterized by less than full deterrence of bogus applications, i.e.,* $p^* < 1$.

Proof: We prove the existence and uniqueness of the stated equilibrium candidate. The agency's best-response function, $e(p)$, is obtained from (6). Firms' best response to examination effort e is

$$d(\theta) = \begin{cases} I & \text{for } \theta < \ell(e,F) \\ B & \text{for } \ell(e,F) \le \theta < h(e) \\ R & \text{for } \theta \ge h(e), \end{cases}$$

leading to a probability $p(e)$ as stated in (7), bounded below by $p_0 \equiv \frac{1-G(h(0))}{1-G(\ell(0,F))} > 0$ and bounded above by 1. For a given F, the upper bound is reached at $e_{\max}$ defined by $h(e_{\max}) = \ell(e_{\max}, F)$. Since $\mathrm{d}h/\mathrm{d}e < 0$ and $\partial\ell/\partial e > 0$, $p(e)$ is monotone increasing in e.

Suppose $\alpha\, L + (1-\alpha)(t_B - t_\varnothing) > 0$ (which will be true at the optimum). Then, the agency's best response is monotone decreasing in p, bounded below by 0 and above by $e_0 \equiv (\gamma')^{-1}(\alpha\, L + (1-\alpha)(t_B - t_\varnothing))$. We conclude that there always exists a unique equilibrium. It involves both some bogus applications and some examination effort unless F deters all bogus applicants even if firms expect $e = 0$. ■

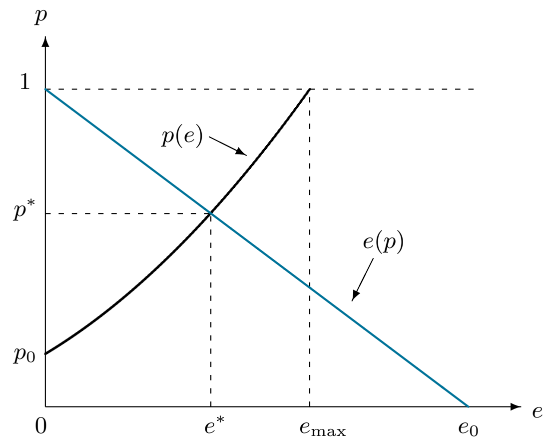

Figure 3: The equilibrium in the hard-information case

Figure 3 depicts the best-response functions for the examination game. The planner's choices of application fee and incentive scheme influence the equilibrium of the game in two ways. The application fee, F, affects the firms' best response: increasing F shifts up the $p(e)$ curve. Thus F also determines the maximum amount of effort for which p is below 1, denoted $e_{\max}$. Specifically, $e_{\max}$ is decreasing in F; we have $e_{\max} = \min\{1 - \pi_R(0)/\pi_B(0), 1\}$ for $F = \max\{\pi_R(0), 0\}$ and $e_{\max} = 0$ for $F = \pi_R(h(0))$. The difference in transfers, $t_B - t_\varnothing$, affects the slope of the agency's best response function: the greater $t_B - t_\varnothing$, the larger (in absolute value) the slope. Summarizing the comparative statics,

- increasing $t_B - t_\varnothing$ raises both e^* and p^*, i.e., examination effort increases and the proportion of bogus applications decreases. It follows that the number of invalid patents issued (given by $(1 - p^*)(1 - e^*)$) is reduced;
- increasing F reduces e^* and increases p^*; the overall effect on invalid patents issued is ambiguous.

One notable result of this analysis is that, to give the agency incentives to provide effort, the planner should pay it for coming up with prior art that constitutes grounds for rejection ($\sigma = B$). Simply put, the planner should reward the agency for *rejecting* applications.

The maximum amount of effort that can be elicited for a given F is $e_{\max}$. At $e_{\max}$, by definition $p = 1$, so no bogus applications are filed. The equilibrium of the game is closer to $e_{\max}$ the larger $t_B - t_\varnothing$ is. This means that, by setting $t_B - t_\varnothing$ sufficiently large, the planner can achieve an outcome that is arbitrarily close to full deterrence. She can use the fee to adjust $e_{\max}$ to the desired level. Hence the following proposition:

Proposition 3. *Suppose $\sigma = B$ is hard information. By choosing $F = F^o$ and $t_B - t_\varnothing$ sufficiently large, the government can achieve an arbitrarily close approximation of the full-commitment outcome: for any ε there exists $t_B - t_\varnothing$ such that $e^* = e^o - \varepsilon$.*

Proof: This is immediate from the discussion above. ■

According to Proposition 3, when the signal is verifiable the planner can achieve (almost) the same outcome as under full commitment. Interestingly, the outcome is unaffected by the agency's degree of intrinsic motivation as measured by α.

4.2 Soft information

In this section, we assume that the signal $\sigma = B$ is soft information. This means that the agency will be tempted to reject even those applications for which it has not discovered any defeating prior art in order to avoid deadweight loss. The incentive scheme must ensure truthful revelation. Technically, the problem becomes one of moral hazard followed by adverse selection: the agency's (unobservable) effort determines the distribution of "types" (in this case, the distribution of signals). We can work backwards from the adverse-selection stage and invoke the revelation principle, according to which a direct revelation mechanism is without loss of generality. Thus, the planner asks the agency to report its signal σ. She pays t_B and rejects if the report is $\sigma = B$. She pays $t_\varnothing$ and grants if the report is $\sigma = \varnothing$.[15]

Consider the case where the agency has exerted equilibrium effort $e^{**} > 0$ and come up with signal $\sigma = B$. For the agency to prefer rejecting to granting a patent, it must be the case that

$$(1-\alpha)t_B \geq (1-\alpha)t_\varnothing - \alpha\, L. \tag{9}$$

If, on the other hand, the agency obtains no signal ($\sigma = \varnothing$), it will prefer granting a patent to refusing it if and only if

$$(1-\alpha)t_B \leq (1-\alpha)t_\varnothing - \alpha[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L], \tag{10}$$

where $\hat{p}(e)$ is the agency's posterior belief that the application is valid given that no evidence to the contrary is found after examination effort e. A final constraint on transfers comes from the possibility of simultaneous deviation: the agency may deviate from both the equilibrium level of effort and truthful reporting, choosing $e = 0$ and nevertheless reporting $\sigma = B$. To

[15] The fact that the mechanism restricts the probability of granting a patent to be 1 when the report is $\varnothing$ and 0 when it is B is without loss of generality. Increasing the probability of a grant after report B above zero, for example, relaxes the adverse-selection constraint, but it also weakens the incentive to provide effort. It turns out that the second effect dominates the first; see the proof of Proposition 4.

rule this out, the equilibrium utility with truthful reporting must be larger than the expected utility with zero effort and report B, or

$$(1-\alpha)t_B \leq (1-\alpha)\big[[p+(1-p)(1-e^{**})]t_{\varnothing}+(1-p)e^{**}t_B\big]-\alpha[pD+(1-p)(1-e^{**})L]-\gamma(e^{**}). \quad (11)$$

Without making further assumptions on the functional form of γ, we cannot say much about how (11) relates to the other constraints. A sufficient condition for (11) to hold, however, is

$$(1-\alpha)t_B \leq (1-\alpha)t_{\varnothing} - \alpha[pD + (1-p)L], \quad (12)$$

as we show in the proof of Proposition 4. In what follows, we replace constraint (11) by (12).[16]

Since by Bayes' rule $\hat{p}(e) = p/[p + (1-p)(1-e)] > p$ for any $e > 0$ and, by assumption, $D < L$, (12) implies (10). Combining conditions (9) and (12), we obtain

$$pD + (1-p)L \leq \frac{1-\alpha}{\alpha}(t_{\varnothing} - t_B) \leq L. \quad (13)$$

To satisfy this condition, transfers must be such that $t_{\varnothing} \geq t_B$, i.e., the patent office must be rewarded for *granting* patents. This contrasts with the hard-information case where the patent office is rewarded for rejecting applications.

We now turn to the incentives to exert effort. Suppose the agency anticipates that it will reject applications if and only if it comes up with signal $\sigma = B$. Then, its optimal level of effort given a prior p that an application is valid is again determined by

$$(1-p)[\alpha L + (1-\alpha)(t_B - t_{\varnothing})] = \gamma'(e).$$

Since for any $\alpha > 0$, $t_B < t_{\varnothing}$ by (13), an incentive scheme that satisfies the condition for truthful grant/refusal decisions tends to reduce effort below the level that would prevail in the hard-information case even in the absence of extrinsic incentives. Note, however, that in the soft-information case, an incentive scheme that doesn't induce truthfulness leads to zero effort.

To maximize examination effort, the planner chooses $t_{\varnothing} - t_B$ at the smallest value consistent with (13), that is,

$$t_{\varnothing} - t_B = \frac{\alpha}{1-\alpha}[pD + (1-p)L]. \quad (14)$$

The planner's hands are tied as to the choice of transfers. She can no longer influence the $e(p)$ function. The only remaining instrument is F, which affects the $p(e)$ curve. But this means that the planner needs to trade off examination effort against the number of bogus applications.

[16] This simplifies the analysis without changing the qualitative results.

Proposition 4. *Suppose Assumption 2 holds. With soft information, the planner is constrained to setting* $(t_\varnothing, t_B)$ *as specified in* (14)*. The equilibrium of the game is such that*

$$\gamma'(e^{**}) = \frac{[1 - G(h(e^{**}))][G(h(e^{**})) - G(\ell(e^{**}, F))]}{[1 - G(\ell(e^{**}, F))]^2}\alpha(L - D).$$

Full deterrence can only be achieved at the cost of no examination effort being made.

Proof: We start by showing that to maximize the agency's incentives, grant probabilities after reports B and $\varnothing$ should be equal to 0 and 1, respectively, as implicitly assumed in the text. Let x_σ denote the grant probability after report σ. The incentive-compatibility constraints (9) and (12) can then be written in a more general way as

$$(1 - \alpha)(t_\varnothing - t_B) \leq \alpha(x_\varnothing - x_B)L \quad (15)$$

$$(1 - \alpha)(t_\varnothing - t_B) \geq \alpha(x_\varnothing - x_B)[pD + (1 - p)L]. \quad (16)$$

The incentives to exert effort are determined by the first-order condition

$$(1 - p)[\alpha L(x_\varnothing - x_B) + (1 - \alpha)(t_B - t_\varnothing)] = \gamma'(e). \quad (17)$$

Maximizing incentives to exert effort is equivalent to maximizing the left-hand side of (17) subject to the constraints (15) and (16). Since incentives increase with $t_B - t_\varnothing$, the binding constraint is (16) and we can replace $(1 - \alpha)(t_B - t_\varnothing)$ by $\alpha[pD + (1 - p)L](x_B - x_\varnothing)$ in (17) to obtain

$$p(1 - p)\alpha(L - D)(x_\varnothing - x_B) = \gamma'(e).$$

Under Assumption 2, the left-hand side of this expression is increasing in $x_\varnothing - x_B$, so incentives are maximized for $x_\varnothing = 1$ and $x_B = 0$.

Next, we establish that (12) is sufficient for (11) to hold. Rewrite (11) as

$$(1 - \alpha)[p + (1 - p)(1 - e^{**})](t_\varnothing - t_B) - \alpha[pD + (1 - p)(1 - e^{**})L] \geq \gamma(e^{**}).$$

Since e^{**} satisfies (6) and γ is convex,

$$(1 - p)e^{**}[\alpha L + (1 - \alpha)(t_B - t_\varnothing)] > \gamma(e^{**}).$$

Thus, a sufficient condition for (11) is

$$(1-\alpha)[p+(1-p)(1-e^{**})](t_\varnothing - t_B) - \alpha[pD+(1-p)(1-e^{**})L] \geq (1-p)e^{**}[\alpha L+(1-\alpha)(t_B - t_\varnothing)]$$

which can be simplified to (12).

Plugging (14) into (6) yields $(1-p)p\alpha(L-D) = \gamma'(e)$, which determines the agency's best response $e(p)$. Firms' best response is still given by $p(e)$ from (7). Combining $e(p)$ and $p(e)$,

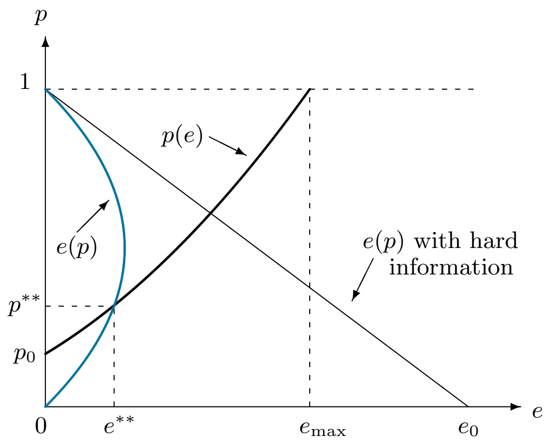

Figure 4: The equilibrium in the soft-information case

one obtains the claimed result. Existence of equilibrium is guaranteed by $p(0) > 0$ which is true by Assumption 3 (insuring finiteness of $h(0)$). While uniqueness is not guaranteed and depends on $G(\cdot)$, in case of multiplicity of equilibria all equilibria are such that e^{**} and p^{**} are strictly less than e^* and p^* since $(1-p)p\alpha(L-D) < (1-p)[\alpha L + (1-\alpha)(t_B - t_\varnothing)]$ for any pair of transfers verifying $t_B \geq t_\varnothing$. ■

According to Proposition 4, when the signal is unverifiable, it is no longer possible to achieve an outcome with both full deterrence of bogus applications and a strictly positive level of examination effort. The equilibrium of the game is illustrated in figure 4. Since t_B and $t_\varnothing$ are determined by (14), the comparative statics reduces to studying the effect of changes in F. As F increases, the $p(e)$ curve again shifts upwards, but now the effects on equilibrium depend on whether one is in the part of the $e(p)$ curve sloping upward or downward. In the upward-sloping part (p between 0 and 1/2), raising F leads to increases in both p^{**} and e^{**}. In the downward-sloping part (p between 1/2 and 1), the comparative statics is the same as before: raising F leads to an increase in p^{**} and a decrease in e^{**}.

It follows that it can never be optimal for the planner to choose F such that the equilibrium is in the upward-sloping part of the $e(p)$ curve; increasing F up to the level where the equilibrium is at the peak of the $e(p)$ curve is unambiguously welfare enhancing. Beyond this point, however, the planner faces a tradeoff: on the one hand, a higher application fee entails fewer bogus applications. On the other hand, the resulting decrease in equilibrium effort reduces the level of innovation. Unlike in the case of hard information, bad patents are inevitable unless the planner sets the fee so high that even in the absence of any examination effort, only true inventors apply for patents. The planner has to choose the lesser of two evils: a situation where no examination takes place ($e = 0$) and bogus applications are deterred through prohibitively large application fees, or a situation with more research but at the expense of some impostors submitting applications *and* a fraction of them obtaining patent protection on their alleged inventions.

A second important difference to the hard-information case is that intrinsic motivation now matters. With soft information, an agency motivated solely by monetary transfers ($\alpha = 0$) cannot be induced to exert any examination effort: it will report whatever pays best. Only if $\alpha > 0$ is a positive level of equilibrium effort sustainable. The higher α, the greater the level of effort that can be sustained for any given F. Thus, the planner should strive to recruit an agency which cares about welfare. This mirrors the result obtained by Iossa and Legros (2004) in the context of auditing with soft information. They find that an auditor must be given property rights in the asset he audits in order for him to exert any auditing effort.

4.3 Discussion

Hard vs soft information

Given how strongly the results of the hard and soft-information cases diverge, it is natural to wonder which is the better model. The assumption that evidence is verifiable is standard in the law and economics literature. But as acknowledged by, e.g., Shin (1998), it may not be a good description of situations involving complex scientific evidence. Patent applications are inherently technical and have increased in complexity over time. Moreover, patentability criteria, and the non-obviousness standard in particular, are often vague, somewhat ill-defined concepts. As noted by Jaffe and Lerner (2004, p. 172), "there is an essentially irreducible aspect of judgment in determining if an invention is truly new. After all, even young Albert Einstein faced challenges while assessing applications (...) in the Swiss Patent Office." In an experiment carried out by the UK Patent Office in 2005, workshop participants were asked to evaluate whether a number of fictitious inventions satisfied different definitions of a "technical contribution" (Friebel *et al.*, 2006).[17] There was large disagreement among participants as to the conformity of the fictitious applications with any given definition. Because of ambiguity in patentability criteria and the technical complexity of applications, patent examiners are likely to have considerable discretion over the decision to grant or reject an application.

Neither hard nor soft information is an accurate description of reality. They are polar cases, and reality is probably somewhere in between. When talking about the implications of the model, all the theoretical results should therefore be taken with a grain of salt. It is nevertheless interesting to note that the soft-information model delivers results which seem

[17] The notion of "technical contribution" was part of a proposed EU directive to deal with software patents; see http://eur-lex.europa.eu/LexUriServ/site/en/com/2002/com2002_0092en01.pdf.

more in line with what we observe empirically. The model can provide a rationale for three observations in particular:

- the apparent laxity of patent examination and the resulting incidence of bad patents. To be sure, if researchers sometimes inadvertently rediscover existing technologies, we will observe bad patents even in the full-commitment case. It is hard to believe, however, that none of the firms filing invalid applications are aware of the relevant prior art. To take the famous example of the patent on the "peanut butter and jelly sandwich"[18] (Jaffe and Lerner, 2004), did the patent holder (jam and jelly maker J.M. Smucker Co.) really think that they had invented this sandwich?

- the controversial compensation scheme in use at the USPTO that rewards examiners for granting patents. Examiners are paid a bonus for achieving certain production targets, where production is measured by the number of applications treated. But a rejection is on average much more time consuming than a grant (for example, applicants can file so-called continuation applications after an initial rejection), so that the production targets basically translate into rewards for grants (Merges, 1999).[19] It is difficult to argue that these rules are not by design: in other, similar settings, such as the refereeing process for academic journals, it is often much easier to reject than to accept candidates. Why make it so difficult for patent examiners to issue a final rejection? And why send a continuation application back to the original examiner, rather than to a different one?

- the high degrees of intrinsic motivation displayed by EPO examiners, as documented in a survey by Friebel *et al.* (2006). This is consistent with the theoretical insight from the model that intrinsic motivation plays a crucial role in the provision of incentives.

None of these observations can be explained by the hard-information model, which predicts that bogus applications are completely deterred, that examiners are rewarded for rejections, and that intrinsic motivation does not matter.

The social costs of good and bad patents

Since the parameters D and L are crucial for the results in the soft-information case, we comment briefly on their interpretation. A difference between the social surplus an invention generates with and without patent protection, as captured by the parameter D, can arise for several reasons. First, monopoly pricing over the lifetime of the patent will lead to (static) deadweight loss. Second and arguably more importantly, in the case of sequential innovation,

[18] United States Patent No. 6,004,596, for a "sealed crustless sandwich."

[19] No such bonus scheme is in use at the European Patent Office (Friebel *et al.*, 2006).

patent protection on first-generation innovations can be a disincentive to later-generation innovators. Third, as an alternative to secrecy, patents encourage disclosure of knowledge and avoid duplication of R&D (Denicolò and Franzoni, 2004a,b). Disclosure may, in a sequential-innovation context, spawn second-generation innovations that would have been impossible had the first-generation invention not been disclosed to the public. Since the third effect is of opposite sign with respect to the first two, it is not a priori clear whether patents are ex post welfare-enhancing or reducing. We have assumed $D \geq 0$ – patents are socially costly ex post – which is probably the predominant view in economics. Section 5.2 below looks at the alternative case where $D < 0$.

The social cost of bad patents, L, also stems partly from monopoly deadweight losses and disincentives to future innovators. Another component of L is the cost of litigation when the patent is challenged in court. Bad patents do not present the offsetting benefit from disclosure, however (since, by assumption, the technologies covered by these patents are not new and cannot be protected by secrecy).

A comparison of L and D is made difficult by the fact that bad patents tend to be challenged and invalidated more often than good patents, which might limit their adverse effects. Still, the vast majority of patent suits are settled out of court, and as Farrell and Shapiro (forthcoming) show, even weak patents (i.e., patents which have a low probability of being valid) can command significant rents. In brief, while there are many arguments both for and against patents, bogus patents present all of the drawbacks but none of the benefits of patents; hence the assumption that $L > D$. Note that this assumption is crucial for the patent office ever to truthfully reveal its signal when information is soft.

5 Robustness and extensions

In this section, we first examine departures from the basic model and check whether the main results remain valid. We then develop a more structural model of the innovation process and investigate under which conditions the profit and welfare functions it generates satisfy Assumptions 2 and 3. Finally, we build on the structural model to look at the effects of expanding the planner's set of instruments to include grant probabilities and differentiated fees.

5.1 Costly public funds

Two of the results from the basic model may appear to rely on public funds being costless: the full-deterrence result of Proposition 1, and Proposition 3 according to which, in the hard-information case, the planner can achieve an arbitrarily close approximation of the optimum by paying a sufficiently large reward. If there is a shadow cost of funds, application fees are a way to raise revenue and reduce the tax burden. Conceivably, this might make it attractive not to deter all bogus applications. A large reward means the examiner obtains a rent, which is socially costly in the presence of a shadow cost of public funds. We now introduce a cost of funds $\lambda > 0$ and show that under plausible conditions, the essence of these results remains intact.

Optimal patent policy

Suppose every dollar of funds that needs to be raised through the tax system costs society $\$(1+\lambda)$ because of tax distortions. Then, the social-welfare function from (3) becomes

$$\int_{\hat{\theta}}^{\infty} W(\theta)\mathrm{d}G(\theta) - (1-e)L[G(\hat{\theta}) - G(\underline{\theta})] - \gamma(e)[1-G(\underline{\theta})] - \lambda\big[t - F[1-G(\underline{\theta})]\big], \qquad (18)$$

where t is the transfer given to the agency and $F[1-G(\underline{\theta})]$ is the revenue from application fees. We can no longer ignore the agency's individual-rationality (IR) constraint. Normalizing the agency's outside opportunity to zero and assuming $\alpha = 0$,[20] IR writes

$$t \geq \gamma(e)[1-G(\underline{\theta})]. \qquad (19)$$

Thus, the planner chooses (e, F, t) to maximize (18) subject to (1), (2), (19) and $\underline{\theta} \leq \hat{\theta}$. Clearly, the IR constraint will be binding. Letting μ denote the multiplier associated with the constraint $\underline{\theta} \leq \hat{\theta}$, and using the fact that $F = (1-e)\pi_B(\underline{\theta})$, the Lagrangian of the problem is

$$\mathcal{L} = \int_{\hat{\theta}}^{\infty} [W(\theta) + (1-e)L]\mathrm{d}G(\theta) - \big(1-G(\underline{\theta})\big)\big[(1+\lambda)\gamma(e) + (1-e)[L - \lambda\pi_B(\underline{\theta})]\big] - \mu(\underline{\theta} - \hat{\theta}).$$

Denote by (e^λ, F^λ) the maximizer of $\mathcal{L}$.

Proposition 5. *Suppose the shadow cost of public funds is λ. A sufficient condition for the optimal patent policy to involve full deterrence is $L/\pi_B(h(0)) > \lambda$. Suppose moreover that λ is not too different from zero. Then, $e^\lambda < e^o$ and $F^\lambda > F^o$.*

Proof: Differentiating the Lagrangian with respect to F, we have

$$\frac{\partial \ell}{\partial F}\Big[g(\underline{\theta})\big[(1+\lambda)\gamma(e) + (1-e)[L - \lambda\pi_B(\underline{\theta})]\big] + \big(1-G(\underline{\theta})\big)(1-e)\lambda\pi_B'(\underline{\theta}) - \mu\Big] = 0.$$

Since $\partial\ell/\partial F > 0$ by Assumption 3, this simplifies to

$$g(\underline{\theta})\big[(1+\lambda)\gamma(e) + (1-e)[L - \lambda\pi_B(\underline{\theta})]\big] + \big(1-G(\underline{\theta})\big)(1-e)\lambda\pi_B'(\underline{\theta}) = \mu.$$

[20] If $\alpha > 0$, IR is easier to satisfy. So there is no additional insight, and computations are more complicated.

If $L/\pi(h(0)) > \lambda$, the left-hand side is strictly positive, so $\mu > 0$, proving the first claim.

Using the fact that $\underline{\theta} = \hat{\theta}$, the derivative with respect to e is

$$-h'(e)W(\hat{\theta})g(\hat{\theta}) - (1 - G(\hat{\theta}))\left[(1+\lambda)\gamma'(e) - \lambda\pi_R' h'(e)\right] + h'(e)g(\hat{\theta})[(1+\lambda)\gamma(e) - \lambda\pi_R] = 0.$$

Rearranging, we obtain

$$-h'(e)g(\hat{\theta})[W(\hat{\theta}) - \gamma(e)] - \gamma'(e)[1 - G(\hat{\theta})] = \frac{\lambda}{1+\lambda}h'(e)\left[g(\hat{\theta})[\pi_R(\hat{\theta}) - W(\hat{\theta})] - \pi_R'(\hat{\theta})(1 - G(\hat{\theta}))\right].$$

The left-hand side is the first-order condition for the case $\lambda = 0$ (equation (4)), which we have shown to be decreasing at e^o (see the proof of Proposition 1) and, by continuity, in its vicinity. The right-hand side is strictly positive since $W(\hat{\theta}) \geq \pi_R(\hat{\theta})$ (social returns exceed private returns) and $h'(e) < 0$. Thus, $e^\lambda < e^o$ if λ is small (so that the intersection of left- and right-hand side is in the vicinity of e^o), and $F^\lambda = \pi_R(h(e^\lambda)) > F^o$. ■

The condition $L/\pi_B(h(0)) > \lambda$ has an intuitive interpretation. The left-hand side represents the (minimum) deadweight-loss-to-profit ratio of bad patents. Since π_B is the maximum amount an impostor is willing to pay for a patent, this measures the cost of raising an additional dollar through the patent system. The right-hand side represents the cost of raising a dollar through the tax system. In general, we should expect that the tax system creates fewer distortions than a patent since taxes are spread over many different markets whereas a patent affects a single market. Thus, the condition seems plausible.

Incentive schemes with multiple applications per examiner

We have derived the result that optimal effort provision requires large rewards under the simplifying assumption that each examiner handles only a single patent application. We now relax this assumption. Instead, we consider the more general case where each examiner treats many applications. We do assume, however, that the number of applications handled per examiner is not large enough for the law of large numbers (LLN) to apply. This seems to be consistent with the workload of real-world examiners.[21] With multiple applications, incentive schemes can be more complex. In particular, they can condition on the total number of grants and rejections. Denote the number of applications handled by a representative examiner by a and the number of rejections by b; the number of grants is thus given by $a - b$. Moreover, assume that the number of examiners is fixed so that a is perfectly informative about the total number of applications received by the patent office, and normalize the examiner's outside opportunity to zero.

[21] According to van Pottelsberghe and François (forthcoming), 96.9 applications per examiner were filed in 2003 at the USPTO, compared to 34.6 at the EPO.

Hard information. We start again with the case where the signal is hard information. Consider the following simple scheme: pay the examiner a bonus of T if the number of rejections (i.e., bad signals) he produces exceeds some threshold $\hat{b}$, and zero otherwise. That is,

$$t(b) = \begin{cases} 0 & \text{for } b < \hat{b} \\ T & \text{for } b \geq \hat{b} \end{cases} \tag{20}$$

Both $\hat{b}$ and T can depend on a; we do not write them as a function of a merely to avoid cumbersome notation. Two questions arise: first, taking the number of applications as given, can $(\hat{b}, T)$ be chosen in such a way that the examiner (a) exerts the desired level of effort and (b) doesn't obtain any rent? Second, can we deal with the problem of multiple equilibria? To see the relevance of the second point, suppose the incentive scheme stipulates a low target $\hat{b}$. If the number of bogus applications is small, the examiner will have to exert considerable effort to reach even such a low target, and there exists an equilibrium with few bogus applications and high effort. There is, however, a second equilibrium where many bogus applications are filed and the examiner provides little effort.

The second problem can be dealt with by making bonus and rejection target depend on the number of applications – provided that the total number of applications gives a good idea of the proportion of bad applications (received by the patent office as a whole, not necessarily by the individual examiner). Formally, this requires imposing a consistency condition: in deriving the examiner's best-response function, we should consider only combinations of p and a that are actually firms' best responses to *some* level of effort e. (This makes sense if all firms hold the same believes about e (not only in equilibrium).) Then the share of good applicants (p) is a deterministic function of the total number of applicants (a), so that we can write it as $p(a)$.

When a is large, the share of bad applicants will be large, too (i.e., $p(a)$ is decreasing in a). To see this, note that $a = (1 - G(\underline{\theta}))/n$ where n is the number of examiners. A large a (i.e., a small $\underline{\theta} = \ell(e, F)$) means that firms expect a low level of effort. This, in turn, leads to a large $\hat{\theta}$, so $p = [1 - G(\hat{\theta})]/[1 - G(\underline{\theta})]$ will be small. Denote by $\bar{a}$ the number of applications such that $p(\bar{a}) = 1$, i.e., when firms expect a level of effort that makes it unprofitable to submit a bogus application for any firm (given F).

In the hard-information case, the number of rejections b coincides with the number of applications for which the examiner finds invalidating prior art ($\sigma = B$). For convenience, we assume that for any $a > \bar{a}$, the random variable b is distributed on the support $[0, a]$ according to a continuous distribution $Q(b; e, a)$, with density $q(\cdot)$.[22]

[22] Formally a continuous distribution conflicts with the assumption that the LLN does not apply because there is an infinity of points in any interval. We have adopted this assumption merely to avoid integer problems;

Assumption 4. *The distribution of b has the following properties:*

(i) $q(b; e, a)$ *is twice differentiable in e and a,*

(ii) $\dfrac{q_e}{q}(b; e, a)$ *is increasing in b,*

(iii) $Q_{ee} \geq 0$,

(iv) $\dfrac{q_e}{q}(a; e^o, a) > \dfrac{\gamma'(e^o)}{\gamma(e^o)}$.

(Subscripts denote partial derivatives.) Property (i) ensures the existence of a solution. Property (ii), known as the monotone likelihood ratio property (MLRP), makes sure that a greater b is indicative of greater effort. Note that MLRP implies first-order stochastic dominance with respect to e. Thus, the larger e, the larger the probability weight that the distribution places on high values of b. This seems natural: the greater e, the higher the probability of finding the bad applications among a given total of a. Property (iii), known as the convexity of the distribution function condition, means that higher effort increases the probability of a large b at a decreasing rate. Finally, property (iv), based on Kim (1997), makes sure that within the feasible set $[0, a]$ there exists a target $\hat{b}$ that permits implementation of a level of effort close to the full-commitment level through incentive scheme (20).

The examiner facing incentive scheme $t(b)$ from (20) solves the following problem:

$$\max_{e} \, (1-\alpha)[1 - Q(\hat{b}; e, a)]T - \alpha L E[y|e, a] - a\gamma(e)$$

where y is the number of bad applications the examiner fails to find. For simplicity, we will focus on the case where $\alpha = 0$ so that the examiner cares only about transfers; none of the qualitative insights depend on this. The first-order condition then is

$$-Q_e(\hat{b}; e, a)T = a\gamma'(e).$$

Thanks to property (iii) in Assumption 4, the first-order condition is necessary as well as sufficient.

Now suppose the planner wants to induce a level of effort $e^o_\varepsilon \equiv e^o - \varepsilon$, where $\varepsilon > 0$ but small, regardless of the number of applications received. Then, the optimal incentive scheme $(\hat{b}, T)$, implementing e^o_ε and leaving no rent to the examiner, must satisfy the following pair of equations for any $a > \bar{a}$:

$$[1 - Q(\hat{b}; e^o_\varepsilon, a)]T = a\gamma(e^o_\varepsilon) \qquad (21)$$

$$-Q_e(\hat{b}; e^o_\varepsilon, a)T = a\gamma'(e^o_\varepsilon) \qquad (22)$$

as long as the number of applications per examiner is not too small, the results should be transposable to the discrete case.

Equation (21) is the examiner's binding individual-rationality constraint, while (22) is the incentive-compatibility constraint.

Proposition 6 (Kim, 1997)**.** *Suppose Assumption 4 holds and that the signal is hard information. With a applications per examiner, the government can achieve an outcome arbitrarily close to full commitment without leaving any rent to the examiner by designing an incentive scheme* $t(b)$ *from* (20) *such that* $(\hat{b}, T)$ *solve equations* (21) *and* (22)*.*

Proof: If the incentive scheme induces examination effort e^o_ε for all a, then firms should rationally expect their application to be screened with that intensity. In equilibrium, $\underline{\theta} = \ell(e^o_\varepsilon, F)$ and $\hat{\theta} = h(e^o_\varepsilon)$. Thus, by choosing $F = F^o$, the government can approximate the full-commitment outcome. What remains to be shown is that the incentive scheme indeed accomplishes its task of inducing e^o_ε. The proof is largely based on Kim (1997); see the proofs of Lemmas 1 and 2 therein.

The scheme in (20) clearly satisfies the limited-liability constraint since it never calls for a negative transfer. What we need to prove is that, for any $a > \bar{a}$, there exist $(\hat{b}, T)$ solving (21) and (22). From (21), $T = a\gamma(e^o_\varepsilon)/[1 - Q(\hat{b}; e^o_\varepsilon, a)]$. Plugging this into (22) and rearranging, we have

$$-\frac{Q_e(\hat{b}; e^o_\varepsilon, a)}{1 - Q(\hat{b}; e^o_\varepsilon, a)} = \frac{\gamma'(e^o_\varepsilon)}{\gamma(e^o_\varepsilon)}. \tag{23}$$

Under MLRP, the left-hand side is nondecreasing in $\hat{b}$ by Lemma 1 in Kim (1997). It is zero at its lower bound, $\hat{b} = 0$ since $Q_e(0; e, a) = 0$ (e has no impact on Q at 0 because $Q(0; e, a) = 0$ for all e). Lemma 2 in Kim (1997) shows that $-Q_e/(1 - Q)$ tends to q_e/q as $b \to a$ by l'Hospital's rule. By property (iv) in Assumption 4, $q_e/q(a; e^o_\varepsilon, a) > \gamma'(e^o_\varepsilon)/\gamma(e^o_\varepsilon)$ when ε is small. Hence, an $\hat{b}$ solving (23) exists. ■

The keys to this result are the following. For the case of a risk-neutral agent protected by limited liability, Kim (1997) has derived the condition for a first-best contract to exist (property (iv) in Assumption 4) and shown that, if any such contract exists, there exists a bonus scheme that implements the first-best allocation. In the proof of Proposition 6, the appropriate target $\hat{b}$ for such a bonus scheme is derived. The optimal target is such that the rate of change in the examiner's expected income equals the rate of change in the cost of effort, both evaluated at e^o_ε. In addition, to rule out undesired equilibria, the rejection target $\hat{b}$ and the bonus T must be adjusted to the number of applications in such a way that e^o_ε is the examiner's best response to *any* number of applications. In that way, firms anticipate that examinations will be examined with intensity e^o_ε regardless of their filing strategy, and there is a unique equilibrium of the game involving only a tiny number of bogus applications.

Proposition 6 demonstrates that the result of Proposition 3 does not hinge on public funds being costless. In fact, with each examiner handling more than one application, attaining the full-commitment outcome does not require leaving a rent to the examiner.[23]

Soft information. In the soft-information case, Proposition 4 shows that the examiner must be paid for accepting applications if he is to reveal his signal truthfully. This implies a tradeoff between incentives to innovate and deterrence of bogus applications. Does this result carry over to the case of multiple applications per examiner? It should be clear that the incentive scheme in (20) does not work when the signal is soft. Regardless of the number of bad signals an examiner has found, he will always claim that the number is greater than $\hat{b}$. This allows him to obtain the bonus and also raises ex post social welfare by avoiding deadweight loss, both of which increase his utility.

Again, we apply the revelation principle and restrict the planner to offer a direct revelation mechanism $\left(\tilde{b}, t(\tilde{b})\right)$, $\tilde{b} \in [0, a]$, that induces the representative examiner to truthfully reveal the number of bad signals he has come up with. The mechanism asks the examiner to designate the applications he has identified as invalid, rejects them and pays a transfer $t(\tilde{b})$ that depends on the number of bad signals reported. Consider an examiner who has exerted the equilibrium level of effort – which we will again denote by e^{**} – and found b bad signals (b is his "type"). His utility from reporting $\tilde{b}$, gross of effort cost $a\gamma(e)$, is given by

$$U(\tilde{b}, r) = \begin{cases} (1-\alpha)t(\tilde{b}) - \alpha\left[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L\right](a-\tilde{b}) & \text{for } \tilde{b} \geq b \\ (1-\alpha)t(\tilde{b}) - \alpha\left([\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L](a-b) + L(b-\tilde{b})\right) & \text{for } \tilde{b} < b \end{cases} \tag{24}$$

where $\hat{p}$ is the examiner's posterior belief that an application is valid given that he has not found any evidence to the contrary. By reporting $\tilde{b}$, the examiner obtains a payoff from the associated transfer, weighted by $1-\alpha$, and a payoff from the resulting level of social welfare, weighted by α. The welfare part is given by $\tilde{b}$ times zero (rejected applicants do not cause any social loss) minus the expected social loss from the $a-\tilde{b}$ accepted applicants. When over-reporting ($\tilde{b} > b$), the examiner's posterior belief about validity is $\hat{p}$ for all of the accepted applicants (he has found no evidence that any of them are invalid), so the expected social loss is $\hat{p}D + (1-\hat{p})L$. When under-reporting, the examiner knows that $b-\tilde{b}$ accepted applications are invalid for sure, while for the remaining $a-b$, his posterior belief about validity is again $\hat{p}$.

Incentive compatibility (IC) requires

$$U(b, b) \geq U(b', b) \quad \forall (b, b') \in [0, a]^2. \tag{25}$$

[23] The assumption that the examiner's outside opportunity is zero is unsubstantial for this result. On the contrary, the results in Kim (1997) suggest that a larger reservation utility makes the existence of a first-best contract more likely.

In particular, (25) must hold for $b' = a$. Reporting a yields the highest possible utility in terms of social welfare (zero). Due to limited liability, type a must nevertheless be given a nonnegative transfer, so $t(a) = 0$ minimizes the incentives to deviate to $\tilde{b} = a$.

As before, we must also account for the possibility of simultaneous deviation: the examiner could shirk ($e = 0$) and report $\tilde{b} = a$. Thus, his expected utility from equilibrium play must be greater than $(1 - \alpha)t(a) = 0$, i.e.,

$$E[U(b, b)|e^{**}] - a\gamma(e^{**}) \geq 0, \tag{26}$$

where the expectation is taken over b. We again resort to the use of a sufficient condition to replace (26), which is that it holds for $e = 0$,[24]

$$\begin{aligned} & E[U(b, b)|0] \geq 0 \\ \iff & t(0) \geq \frac{\alpha}{1 - \alpha}[pD + (1 - p)L]a. \end{aligned} \tag{27}$$

The following proposition characterizes incentive-compatible transfer schemes.

Proposition 7. *Suppose the signal is soft information and each examiner handles a applications. An optimal transfer scheme inducing truthful revelation must satisfy*

(i) $0 \leq t(b) \leq \frac{\alpha}{1 - \alpha}[pD + (1 - p)L]a$

(ii) $-\frac{\alpha}{1 - \alpha}L \leq t'(b) \leq -\frac{\alpha}{1 - \alpha}[\hat{p}(e^{**})D + (1 - \hat{p}(e^{**}))L]$.

Proof: Notice first that the utility function in (24) satisfies single-crossing: the indifference curves it generates are parallel for all r, with a kink at r. Thus, it is sufficient to look at local incentive compatibility. We begin with property (ii). Suppose the first inequality is not satisfied, i.e., $(1 - \alpha)t'(r) < -\alpha L$, even though truthful reporting is optimal for all r. Take any $0 < r \leq a$, and consider a small deviation from truthful reporting, $\tilde{r} = r - \varepsilon$ where $\varepsilon > 0$. Taking a first-order Taylor expansion of $U(\tilde{r}, r)$ around r, we have

$$\begin{aligned} U(r - \varepsilon, r) &= (1 - \alpha)[t(r) - t'(r)\varepsilon] - \alpha\big([\hat{p}(e^{**})D + (1 - \hat{p}(e^{**}))L](a - r) + L(r + \varepsilon)\big) \\ &> (1 - \alpha)t(r) + \alpha L\varepsilon - \alpha\big([\hat{p}(e^{**})D + (1 - \hat{p}(e^{**}))L](a - r) + L(r + \varepsilon)\big) \\ &= U(r, r), \end{aligned}$$

[24] If it is optimal for the examiner to exert positive effort given that he anticipates truthfully reporting b, it must be the case that his expected utility is greater than with $(e = 0, \tilde{b} = 0)$. If his expected utility at $e = 0$ is greater with truthful reporting ($\tilde{b} = 0$) than with reporting $\tilde{b} = a$, it must a fortiori be greater at the equilibrium effort.

contradicting the optimality of reporting r. Similarly, suppose the second inequality is not satisfied, i.e., $(1-\alpha)t'(r) > -\alpha\big[[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L]$, and consider a deviation $\tilde{r} = r + \varepsilon$:

$$\begin{aligned} U(r+\varepsilon, r) &= (1-\alpha)[t(r) + t'(r)\varepsilon] - \alpha[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L](a - r - \varepsilon) \\ &> (1-\alpha)t(r) - \alpha[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L]\varepsilon - \\ &\quad - \alpha[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L](a - r - \varepsilon) \\ &= U(r, r), \end{aligned}$$

again a contradiction.

Thus, we have proved property (ii). Property (i) follows from limited liability ($t(r) \geq 0$), moral-hazard considerations ($t(0)$ must be as small as possible consistent with (26), and thus never greater than $\frac{\alpha}{1-\alpha}[pD + (1-p)L]a$ by (27)), and the fact that, by property (ii), $t(r)$ is strictly decreasing in r. ■

The optimal transfer scheme is again decreasing with b: the examiner is rewarded for granting patents. Note that, if the examiner is motivated purely by monetary transfers ($\alpha = 0$), incentive compatibility is trivially satisfied for $t(b) = 0$ for all b (or, more generally, any transfer that is independent of b). The examiner is simply indifferent between all possible reports, so under standard assumptions, he will report the truth. The real problem lies at the effort-provision stage: with a constant transfer the examiner does not have any incentive to exert effort. Intrinsic motivation ($\alpha > 0$) is a necessary condition for effort provision.

The results from section 4 thus carry over to the case of multiple applications per examiner.[25] Let us point out that the assumption of inapplicability of the LLN is critical for this result. If the LLN does apply, then there is no adverse-selection problem because effort and the number of applications uniquely determine the number of bad signals. In that case, the patent office does not have to be paid for granting patents. Nevertheless, monetary transfers cannot induce it to exert effort. The incentive scheme from the hard-information case, for example, does not work: any target rejection rate set by the planner can be achieved by rejecting at random. Regardless of whether the LLN applies, thus, the full-commitment outcome is out of reach.

5.2 Positive ex post welfare effects of patents

The result in Proposition 4 according to which incentive compatibility requires paying the patent office for granting patents is based on the premise that, from an ex post point of view, patents are costly to society ($D \geq 0$, Assumption 2). While this is probably the standard

[25] This is true qualitatively speaking. Equation (6) determining effort provision will of course have to be adapted to the multiple-application case.

perception in economics, there exist environments where the ex post welfare effect of patents is positive. For example, if firms file for a patent in early stages of their research project, a rejection may lead them to abandon the project. A patent grant is tied to the condition that the inventor describe his invention to the public. Disclosure promotes diffusion, inspires follow-on inventions, and prevents duplication of R&D.[26]

Alternatively, even if $D \geq 0$, patent examiners may for some reason care about something other than (ex post) welfare. They may care about making correct decisions (i.e., approving valid applications, and rejecting invalid ones),[27] or they may take into account the effects of their effort choice on incentives to innovate.[28]

For any of these reasons, the modeling approach we have adopted could be inappropriate. We now consider how an alternative approach – allowing ex post welfare effects to be positive – affects results. Suppose that contrary to Assumption 2, $D < 0$.[29] The analysis of the optimal full-commitment policy and of the hard-information case are unaffected by these changes, so we focus on the soft-information case. Truthful revelation still imposes an upper bound on the difference between transfers, formalized in constraint (12):

$$t_B - t_\varnothing \leq -\frac{\alpha}{1-\alpha}[pD + (1-p)L].$$

Now, however, there are two cases. If $p < L/(L-D)$, the term in square brackets is positive, so nothing changes: upon observing no signal, the agency is still tempted to reject the application and needs to be paid not to. If $p \geq L/(L-D)$, the term in square brackets is negative, so the agency prefers to grant a patent when $\sigma = \varnothing$. It can then be paid for rejecting patents, which improves its incentives for effort provision. But the amount it can be paid is limited. That soft information makes the full-commitment unattainable and creates a tradeoff between innovation and deterrence is a robust result.

Whether an observed incentive scheme will reward for grants or rejections depends on the

[26] See also section 5.3.

[27] This might be interpreted as a rough proxy for career concerns.

[28] The analogy is that of a judge: even though it may often be ex post inefficient to convict a criminal (keeping him in prison is costly for society), the judge is supposed to take into account ex ante considerations. There is probably little reason to worry that an intrinsically motivated judge will acquit a defendant he believes to be guilty to save on costs of imprisonment. It is debatable whether the same argument applies to patent examiners.

[29] This is one alternative approach. As the previous discussion suggests, a second alternative consists in letting the agency care about ex ante welfare, $W(\theta)$. In that case, the agency has to form an expectation over the welfare effect it creates by granting a patent to a firm having done genuine research, given by

$$E[W(\theta)|\hat{\theta}] = \int_{\hat{\theta}}^{\infty} W(\theta)\mathrm{d}G(\theta)/[1-G(\hat{\theta})].$$

Since $\hat{\theta}$ is negatively related to p, this introduces an additional term that depends on p in the agency's best-response function. The analysis becomes slightly more complex, but the qualitative results are the same as with a constant $D < 0$.

equilibrium value of p, which is a function of the parameters α, L, D, the cost function $\gamma(\cdot)$, and the distribution of θ. An incentive scheme is more likely to reward for grants if α is small (the less important intrinsic motivation, the lower the equilibrium p) and D is small in absolute value (the threshold $L/(L-D)$ increases with D, while the equilibrium p decreases). The effect of L is ambiguous: it increases the equilibrium p but also the threshold $L/(L-D)$.

This adds an interesting twist to the argument about empirically observed incentive schemes. It suggests that, as intrinsic motivation increases, it becomes less likely that the incentive scheme rewards examiners for granting patents. If EPO examiners care more about doing high-quality work than examiners at the USPTO, this would explain why USPTO examiners, but not EPO examiners, are rewarded for granting patents (see section 4.3). Unfortunately, while there is evidence of EPO examiners' intrinsic motivation, no such evidence exists for their USPTO counterparts.

5.3 A more structural model

In the basic model we have used reduced-form profit and welfare functions. This section presents a structural model which, under some conditions, leads to profit and welfare functions with the assumed properties.

The economy is populated by a mass 1 of firms and a single consumer with additively separable preferences who consumes one unit of each available good.[30] There is a stock of ideas which are indexed by v, where $v \in [0, \infty)$ is the consumer's valuation for the product that can be derived from the idea. The stock is large, in the sense that there are many ideas that the consumer values v, for any v. Ideas are of two types, old and new. An old idea already exists as a product. Unless protected by a patent, any firm can produce it, i.e., the production technology is known. A new idea must be turned into a product. Turning an idea v into a product requires two things: coming up with the idea, and investing an amount $\psi(v)$ in research. The research phase is about discovering the production technology. Once the technology is discovered, it becomes commonly known, but the innovator can secure a profit by applying for a patent at the patent office. Patent protection is perfect for new ideas, but imperfect for old ideas: if a firm escapes detection by the patent office and manages to patent an old idea, protection fails with probability $1-\beta$, in which case, again, any firm can produce it.[31] The marginal cost of production is 0 for all products. The consumer's utility from consuming a product he values v and buys at price P is $v - P$. A monopolist will thus set $P = v$ and extract the entire surplus, whereas competition will drive the price down to 0.

[30] This setup rules out deadweight losses resulting from monopoly pricing. Below, we do consider the deadweight losses that patents can cause in a sequential-innovation setting, which are arguably more important.

[31] This should be interpreted as a court invalidating the patent.

Firms differ in creativity θ. Creativity matters because the value of an idea is not a priori observable. Whether they are seeking a new idea for research or an old idea for filing a bogus application, more creative firms come up with more valuable ideas. To be specific, a firm θ will find an idea the consumer values $v = \theta$. (This is a normalization: the creativity parameter θ is defined as the value of the idea the firm can produce. The real assumption here is that the relation between creativity and idea is deterministic, rather than stochastic.) In this framework, if granted a patent, a firm having done research makes a profit (gross of patent application fees) of $\pi_R(\theta) = \theta - \psi(\theta)$, while a firm having filed a bogus application makes an expected profit of $\pi_B(\theta) = \beta\theta$.

We now give conditions on $\psi(\cdot)$ under which π_R and π_B satisfy Assumption 3. Since $\pi_B(0) = 0$ and $\pi_R(0) = -\psi(0)$, having $\pi_B(0) > \pi_R(0)$ requires $\psi(0) > 0$. Concavity of $\pi_R(\cdot)$ is equivalent to convexity of $\psi(\cdot)$, so we need $\psi'' \geq 0$. Given convexity of $\psi(\cdot)$, it is sufficient that $\pi'_R > \pi'_B$ hold in the limit as θ tends to ∞ for it to hold at every θ. Thus, we must have $\psi' > 0$ and $\lim_{\theta\to\infty} \psi'(\theta) < 1 - \beta$. Finally, since $\lim_{\theta\to\infty} \pi'_R(\theta) > \beta > 0$, we have $\lim_{\theta\to\infty} \pi_R(\theta) = \infty$.

We now turn to welfare. In the simple model presented so far, neither good nor bad patents cause any (ex post) welfare loss; they merely transfer surplus from the consumer to the patentees. The welfare effect of an innovation is equal to the profit it creates, i.e., $W(\theta) = \pi_R(\theta)$. But consider the following extension of the model that accounts for secrecy and sequential innovation.

Assume that a new idea can be protected either by a patent or by keeping it secret. As in Denicolò and Franzoni (2004a), protection by secrecy is imperfect: with probability $1 - s$, the secret leaks out and becomes commonly known. For the previous analysis to remain valid, it must be the case that all firms prefer patenting over secrecy. This requires imposing an upper bound on the effectiveness of secrecy. We derive this upper bound below; for now, we take it as given that patenting is more profitable than secrecy.

Any idea that is turned into a product and disclosed to the public (i.e., either an old idea or a new idea that is patented or leaked out) inspires a firm (different from the first innovator) to develop an innovation building on the original one. Following Bessen and Maskin (forthcoming), we assume that (a) there is no replacement effect, so the second innovation does not reduce the value of the first one; and (b) the second innovation falls within the breadth of the patent on the first innovation. Practicing the second innovation thus requires a license from the patent holder.

The value of the second innovation is assumed to be independent of the first. Casual empirical evidence supports this assumption: first-generation innovations with little stand-alone value may lead to highly profitable second-generation innovations. Conversely, highly profitable first-generation products do not necessarily lead to significant second-generation products. Suppose moreover that there is asymmetric information on the value of the second innovation. While the second innovator knows that the consumer values his innovation v, the first innovator knows only that v is drawn from a distribution $M(\cdot)$.

Licensing negotiations take place before the second innovator invests $\psi(v)$.[32] Given a license fee ϕ, the second innovator will buy a license if and only if $v - \psi(v) \geq \phi$. Let v_ϕ be the cutoff value defined by $v_\phi - \psi(v_\phi) = \phi$. The patent holder chooses ϕ to maximize his expected licensing revenue given by

$$[1 - M(v_\phi)]\phi.$$

Denote the optimal license fee by ϕ^*. Clearly, $\phi^* > 0$ and $v_{\phi^*} > v_0$, so some second-generation innovators do not find it worthwhile to obtain a license, and their idea is lost. This is how patents can cause deadweight loss in this model.

Let us compare the welfare effects created by an innovator θ when he obtains a patent and when he doesn't. First, define

$$V \equiv \int_{v_0}^{\infty} [v - \psi(v)] \mathrm{d}M(v)$$

as the expected social value of a second-generation innovation and

$$K \equiv \int_{v_0}^{v_{\phi^*}} [v - \psi(v)] \mathrm{d}M(v)$$

as the loss of social value caused by licensing under asymmetric information. The welfare effect of an innovator θ who obtains a patent is

$$W(\theta) = \theta - \psi(\theta) + V - K.$$

If the same innovator is refused a patent, he still has the possibility to protect his innovation through secrecy.[33] The corresponding welfare effect is

$$S(\theta) \equiv \theta - \psi(\theta) + (1 - s)V,$$

since the second innovation can only occur if the first one leaks out. The difference between the two expressions corresponds to the parameter D from the previous sections,

$$D = S(\theta) - W(\theta) = K - sV. \tag{28}$$

[32] The qualitative results are unaffected by this assumption. In fact, it biases the analysis in favor of patents: ex ante bargaining is more efficient than ex post bargaining because of the hold-up problem that arises once the R&D investment is sunk (Green and Scotchmer, 1995).

[33] This is true assuming that applications do not become public unless a patent is granted. We should note that this is inconsistent with current statutes at the European Patent Office which require that all applications be published 18 months after filing. However, even if a rejected application is made public, the innovation may not diffuse as well as under patent protection. The inventor may have some tacit knowledge, and without protection, he has no incentive to seek potential licensees.

Equation (28) highlights the tradeoffs associated with patents after innovation has occurred. On the one hand, patents restrict access to innovative knowledge (K). But on the other hand, they encourage disclosure, leading to subsequent innovation that could not have happened under secrecy ($-sV$). Granting a patent decreases ex post welfare ($D > 0$) if $s < K/V$ and increases ex post welfare ($D < 0$) if $s > K/V$.

For bogus applications, ex post welfare if they do not obtain a patent is V. Since the idea was already known, it cannot be protected by secrecy, and second-generation innovation always occurs. If bogus applicants obtain a patent, they can ask for a license fee – thereby excluding some second-generation innovators – unless their patent is invalidated. Welfare is $\beta(V - K) + (1 - \beta)V = V - \beta K$. The difference in welfare caused by granting a patent is

$$L = \beta K. \tag{29}$$

Unlike patents on true innovations, patents on old ideas can never be welfare enhancing, neither ex post nor ex ante. Whether ex post these bad patents are more or less costly than valid patents depends on parameters, though. From equations (28) and (29), we obtain the following condition for L to be greater than D:

$$s \geq (1 - \beta)\frac{K}{V}.$$

This condition is more likely to be satisfied the more effective secrecy is in protecting innovations, the less likely bad patents are to be overturned, and the smaller the loss of second-generation innovation due to licensing.

Finally, we check whether the modifications that the sequential-innovation model entails affect the conformity of profit functions with Assumption 3; we also derive an upper bound on s for secrecy not to be an option ex ante. Denote the licensing revenue by $\Phi \equiv [1 - M(v_{\phi^*})]\phi^*$. The profit from research now is $\pi_R(\theta) = \theta - \psi(\theta) + \Phi$, while the profit from a bogus patent is $\pi_B(\theta) = \beta\big[\theta + \Phi\big]$. Since in both cases, the additional term is constant with respect to θ, first and second derivatives are unaffected. To have $\pi_R(0) < \pi_B(0)$, we now need $\psi(0) > (1 - \beta)\Phi$.

What is the upper bound on the effectiveness of secrecy to rule out that firms choose secrecy over patenting? Consider a firm which has sunk the R&D investment of $\psi(\theta)$ and needs to decide how to protect its innovation. Patenting yields a profit of $\theta + \Phi - F$ while secrecy yields $s\theta$. The innovator who is worst off with patents is the one at $\hat{\theta}$. Suppose the application fee is maximum, so that this innovator makes 0 profit (net of research costs), i.e. $F = \pi_R(\hat{\theta})$. Thus, his gross profit is equal to the cost of research, $\hat{\theta} + \Phi - F = \psi(\hat{\theta})$. He prefers patenting over secrecy if $s\hat{\theta} < \psi(\hat{\theta})$. For secrecy never to be an option, this must hold for all possible $\hat{\theta}$. But by the convexity of $\psi(\cdot)$, it is sufficient that $\psi'(0) > s$. For this to be possible given Assumption 5, we need $s < 1 - \beta$. Note that $s\theta < \psi(\theta)$ for all θ implies that secrecy is ineffective as an incentive mechanism; the patent system is essential to induce any firm to invest in research. That patents may also encourage rent-seeking is a necessary evil.

We summarize the findings from this section in the following assumption which, in the framework of this more structural model, is equivalent to Assumptions 2 and 3.

Assumption 5. *The cost of research $\psi(\theta)$ and the secrecy parameter s satisfy*

(i) $\psi(0) > (1-\beta)\Phi$

(ii) $\psi'(0) > s$ *and* $\lim_{\theta\to\infty} \psi'(\theta) < 1-\beta$

(iii) $\psi'' \geq 0$

(iv) $(1-\beta)K/V \leq s < \min\{K/V, 1-\beta\}$.

5.4 Expanding the planner's set of instruments

Probabilistic patent grants

We have assumed that the planner uses a deterministic scheme to award patents to applicants: upon receiving a given patentability signal, the patent office issues a given decision (namely, grant after $\sigma = \varnothing$, rejection after $\sigma = B$). Conceivably, the planner may instead want to use a probabilistic scheme. By not always issuing a patent when no signal is found, she avoids deadweight loss.[34] We show here that it is not clear whether the planner would want to use probabilistic grants even within the framework of the model – intuitively, they reduce firms' incentives to do research. We then argue that probabilistic grants are undesirable for several reasons that are outside the model.

Suppose that the planner issues a patent with probability x after the patent office reports no signal. Extending the structural model introduced in the previous section, the thresholds $\hat{\theta}$ and $\underline{\theta}$ are now determined by

$$x[\hat{\theta} + \Phi] - \psi(\hat{\theta}) = (1-e)x\beta[\hat{\theta} + \Phi] \tag{30}$$

defining the implicit function $\hat{\theta} = h(e, x)$, and

$$(1-e)x\beta[\underline{\theta} + \Phi] = F \tag{31}$$

defining the implicit function $\underline{\theta} = \ell(e, F, x)$. The potential welfare generated by an invention being $W(\theta) + D = \theta - \psi(\theta) + V$, the planner's objective becomes

$$\max_{e,F,x} \int_{\hat{\theta}}^{\infty} [\theta - \psi(\theta) + V] \mathrm{d}G(\theta) - x\big[(K - sV)[1 - G(\hat{\theta})] + (1-e)\beta K[G(\hat{\theta}) - G(\underline{\theta})]\big] - \gamma(e)[1 - G(\underline{\theta})]$$

[34] After $\sigma = B$, it is optimal to always reject the application.

subject to $\underline{\theta} \leq \hat{\theta}$, (30), and (31).

Once again, the inequality constraint will be binding.[35] Using this fact, the derivative of the objective with respect to x is

$$-h_x[\hat{\theta} - \psi(\hat{\theta}) + V - (1-x)(K - sV) + \gamma(e)]g(\hat{\theta}) - (K - sV)[1 - G(\hat{\theta})]. \tag{32}$$

By the implicit function theorem,

$$h_x(e, x) = \frac{\hat{\theta}[1 - \beta(1-e)]}{\psi'(\hat{\theta}) - x[1 - \beta(1-e)]}.$$

For values of x close to 0, h_x is positive, while for values close to 1, it is negative by Assumption 5. We are interested in whether decreasing x below 1 is welfare-enhancing. The fact that $h_x(e, 1) < 0$ tells us that decreasing x below 1 locally leads to a reduction in the share of firms doing research. Intuitively, since firms doing research are more likely to obtain a patent than firms filing bogus applications, a decrease in x hurts them more than the impostors. As a result, the relative attractiveness of research declines.

Evaluating (32) at $x = 1$ yields

$$-h_x(e, 1)\left[\hat{\theta} - \psi(\hat{\theta}) + V + \gamma(e)\right]g(\hat{\theta}) - (K - sV)[1 - G(\hat{\theta})].$$

This expression cannot be signed unless $K \leq sV$ (i.e., $D \leq 0$), in which case it is positive.

In general, thus, it is ambiguous whether reducing x below 1 leads to an improvement or a reduction of welfare, and for $D \leq 0$, it certainly leads to a welfare reduction (at least locally). Decreasing the grant probability is also unambiguously bad for the patent office's incentives to exert effort. Under hard information, it reduces the expected welfare impact of not coming up with defeating prior art. Under soft information, there is a potentially beneficial effect: reducing the grant probability relaxes the incentive-compatibility constraint. However, the proof of Proposition 4 shows that the negative effect on the incentive for effort provision dominates.

Apart from the fact that the model makes no clear statement in favor of randomizing the grant decision, there are factors outside the model that tend to make such a policy unattractive. First, there is again a commitment issue: it is difficult for firms to verify whether the planner adheres to announced grant probabilities. This creates a temptation for the planner (or the patent office) to behave opportunistically and grant fewer patents than

[35] The proof is the same as the corresponding one for Proposition 1; the first-order condition with respect to F is the same except that it is multiplied by x, so the multiplier will be strictly positive. Formally, denoting the multiplier by μ, we have

$$\frac{\partial \ell}{\partial F}\left[xg(\underline{\theta})[(1-e)\beta K + \gamma(e)] - \mu\right] = 0$$

which implies $\mu > 0$ (unless $x = 0$, which can be ruled out as a solution).

announced. Since firms will anticipate such opportunism, any positive effects of reducing grant probabilities may be nullified. There is also a more practical issue concerning the legality or constitutionality of randomizing over the grant decision. Finally, the model frames R&D investment as a 0-1 decision. More generally, firms have to decide how much to invest. In a context where the amount of investment is a decision variable, decreasing the grant rate adversely affects both the share of firms doing research and the amount they invest.

Differentiated fees and inadvertent re-invention

In the basic model, we have restricted the planner to charging a uniform application fee for everyone. This rules out an intuitively appealing way of reducing the attractiveness of imposture: charging firms caught submitting bogus applications a fine. The absence of this kind of fine in practice is somewhat of a puzzle. In the framework of the basic model, it would be possible to attain the first best simply by setting a very large fine and exerting small but positive examination effort.[36] The absence of fines in practice could be related to the possibility that even firms doing genuine research may sometimes inadvertently rediscover old ideas. Penalizing them too heavily would reduce incentives to engage in research. This section extends the basic model by allowing the planner to differentiate fees according to the signal reported by the patent office and by introducing inadvertent re-invention.

Suppose a firm doing research comes up with an original invention only with probability ν, while with probability $1-\nu$ it rediscovers an existing one. The planner charges a fee F_σ depending on the signal reported by the patent office. Activity R then procures a type-θ firm an expected payoff of

$$\Pi_R(\theta, e, F_\varnothing, F_B) \equiv [\nu + (1-\nu)(1-e)\beta](\theta+\Phi) - \psi(\theta) - F_\varnothing - (1-\nu)e(F_B - F_\varnothing),$$

while its expected payoff from activity B is

$$\Pi_B(\theta, e, F_\varnothing, F_B) \equiv (1-e)\beta(\theta+\Phi) - F_\varnothing - e(F_B - F_\varnothing).$$

The threshold $\hat{\theta} = h(e, F_\varnothing, F_B)$ is determined by

$$\Pi_R(\hat{\theta}, e, F_\varnothing, F_B) = \Pi_B(\hat{\theta}, e, F_\varnothing, F_B), \tag{33}$$

while $\underline{\theta} = \ell(e, F_\varnothing, F_B)$ is determined by

$$\Pi_B(\hat{\theta}, e, F_\varnothing, F_B) = 0. \tag{34}$$

The condition determining $\hat{\theta}$, (33), can be rewritten as

$$\nu\big[[1-\beta(1-e)](\hat{\theta}+\Phi) + e(F_B - F_\varnothing)\big] - \psi(\hat{\theta}) = 0. \tag{35}$$

[36] This is the classic argument by Becker (1968).

The left-hand side is increasing in e for any positive ν: greater examination effort continues to make research relatively more attractive than imposture. Intuitively, the probability of having a valid invention is strictly greater when doing research.

Even in the absence of bogus applications, it is now ex post efficient to examine applications. Let e^{ν} denote the ex post efficient level of effort, determined by $(1-\nu)\beta K = \gamma'(e^{\nu})$. To accommodate the new setting, we replace Assumption 5 by

Assumption 6. *The cost of research $\psi(\theta)$ and the secrecy parameter s satisfy*

(i) $\psi(0) > \nu(1-\beta(1-e^{\nu}))\Phi$

(ii) $\psi'(0) > s$ *and* $\lim_{\theta\to\infty} \psi'(\theta) < \nu(1-\beta)$.

(iii) $\psi'' \geq 0$

(iv) $s < \nu(1-\beta)$.

Property (i) means that the ex post efficient level of effort does not suffice to make research more profitable than imposture for all firms when fees are zero. This is a minimal assumption for bogus applications to ever occur. Property (ii) is needed for single-crossing. It is slightly stronger than property (ii) in Assumption 5 and ensures that the derivative of the left-hand side of (35) with respect to $\hat{\theta}$, given by $\nu[1-\beta(1-e)] - \psi'(\hat{\theta})$, is positive. It also implies that the social value of research increases faster with θ than its cost. If this were not the case, it would not be socially desirable that high-creativity firms invest in research. (Note that it does not rule out that the private value of research exceeds the social value.)

A first-best allocation in this setting is such that

(i) no firms submit bogus applications, i.e. $\underline{\theta} = \hat{\theta}$ and thus $p = 1$;

(ii) the patent office's examination effort is ex post efficient given $p = 1$, i.e. $e^{FB} = e^{\nu}$;

(iii) all firms that, given e^{FB}, can produce inventions whose expected social value exceeds the cost of research invest, while all others stay idle, i.e. $\hat{\theta} = \theta^{FB}$ defined by $\nu(\theta^{FB} + V - K) - (1-\nu)(1-e^{\nu})\beta K = \psi(\theta^{FB})$.

Proposition 8. *Suppose there is inadvertent re-invention and the planner can differentiate fees according to σ. Suppose also that Assumption 6 holds. The following policy implements the first-best allocation:* $\big(e = e^{FB}, F_{\varnothing} = \theta^{FB} + \Phi - \psi(\theta^{FB})/\nu, F_B = (1-e^{FB})/e^{FB}\,[\psi(\theta^{FB})/\nu - (1-\beta)(\theta^{FB}+\Phi)]\big)$.

Proof: We show first that given the stated policy, $\hat{\theta} = \theta^{FB}$. Substituting for $(e, F_\varnothing, F_B)$, and noting that $F_B - F_\varnothing = [\psi(\theta^{FB})/\nu - (1-\beta(1-e^{FB}))(\theta^{FB}+\Phi)]/e$, we have

$$\nu\big[[1-\beta(1-e^{FB})](\hat{\theta}+\Phi) + \psi(\theta^{FB})/\nu - (1-\beta(1-e^{FB}))(\theta^{FB}+\Phi)\big] = \psi(\hat{\theta})$$

$$\iff \psi(\hat{\theta})/\nu - (1-\beta(1-e^{FB}))\hat{\theta} = \psi(\theta^{FB})/\nu - (1-\beta(1-e^{FB}))\theta^{FB}.$$

By Assumption 6, $\psi(\theta)/\nu - (1-\beta(1-e))\theta$ is strictly monotonic, so the only solution is $\hat{\theta} = \theta^{FB}$.

Second, substituting in (34),

$$(1-e^{FB})\beta(\underline{\theta}+\Phi) - [\theta^{FB}+\Phi-\psi(\theta^{FB})/\nu] - [\psi(\theta^{FB})/\nu - (1-\beta(1-e^{FB}))(\theta^{FB}+\Phi)] = 0,$$

the unique solution of which is $\underline{\theta} = \theta^{FB}$. Finally, having established that $\underline{\theta} = \hat{\theta} = \theta^{FB}$, $e = e^{FB}$ implements the first-best allocation. ■

Proposition 8 says that differentiated fees allow the planner to achieve the first-best outcome even when firms sometimes inadvertently re-invent things. To understand the intuition for this result, let us compute the effect of the fees F_σ on the threshold $\hat{\theta}$. Since only the difference $F_B - F_\varnothing$ matters, we take the derivative

$$\frac{\partial h}{\partial(F_B - F_\varnothing)} = \frac{\nu e}{\psi'(\hat{\theta}) - \nu[1-\beta(1-e)]} < 0,$$

where the inequality follows from Assumption 6. Thus, as long as $e > 0$, raising $F_B - F_\varnothing$ can achieve the same as increasing effort – encourage research – but at no cost. (It also acts as a deterrent.) Since the planner does not have to trade off the benefits of research against the costs of examination, she can achieve the welfare-maximizing level of research.

Whether $F_B > F_\varnothing$ depends on whether θ^{FB} is smaller or greater than $h(e^{FB}, 0, 0)$. If it is smaller, the social returns to research exceed the private incentive to invest, so the planner wants to encourage research and sets $F_B > F_\varnothing$. In the opposite case, the planner wants to discourage research and sets $F_B < F_\varnothing$. To see this, note that from Proposition 8 we get

$$F_B > F_\varnothing \iff \psi(\theta^{FB}) > \nu[1-\beta(1-e)](\theta^{FB}+\Phi).$$

By definition, $\psi(\theta^{FB}) = \nu(\theta^{FB}+V-K) - (1-\nu)(1-e^\nu)\beta K$: at the first-best threshold, the cost equals the social returns to research. The right-hand side of the inequality represents the private incentive to choose research at θ^{FB} (gross of fees; see equation (35)). Saying that the private incentive is smaller than the social returns is equivalent to saying that the cutoff the planner chooses is below the one that firms would choose in the absence of fees, i.e. $\theta^{FB} < h(e^{FB}, 0, 0)$.

The planner uses the level of $F_\varnothing$ to deter bogus applications. Alternatively, if at θ^{FB}, private returns to research are lower than its cost, she sets a negative fee to satisfy firms' participation constraint. From Proposition 8, we have

$$F_\varnothing < 0 \iff \theta^{FB} + \Phi - \psi(\theta^{FB})/\nu < 0.$$

This expression equals firm θ^{FB}'s expected profit from doing research, net of the expected fine, $(1-\nu)e(F_B - F_\varnothing)$. If this firm makes a negative profit, it needs to receive a subsidy to be induced to do research. In that case, all firms that apply for a patent and are not found to be impostors are subsidized.[37]

We have derived these results for the case where the planner directly controls e. We now comment briefly on what happens when examination is delegated to the patent office. With hard information, to the extent that the patent office does not care as much about ex post efficiency as the planner ($\alpha < 1$), the planner only needs to complement the patent office's intrinsic motivation to achieve the first-best. Even more strikingly, soft information does not hurt efficiency much either: even though it is hard to incentivize the patent office through extrinsic rewards, the planner can still achieve full deterrence by increasing the fine F_B.

Given how effectively differentiated fees work in the model, it is puzzling that they are not used in practice. We can only speculate about the cause of their absence. Possibly, political-economy concerns play a role. Fines could be misused by the government. To raise revenues, it might be tempted to reject more than warranted.

5.5 Other issues

Reputation as a remedy to the commitment problem

This paper models patent examination as a one-shot game where the patent office chooses its effort after firms have made R&D investments. One might object that in practice patent examination is a repeated game, raising the question as to whether the patent office can develop a reputation for rigorous screening. Such reputation concerns might bring its ex post and ex ante incentives more in line. Even though individual firms may not be able to observe whether the agency adheres to a policy of rigorous screening, the patent attorneys they charge with prosecuting their applications should be able to get a good idea of the agency's actual policy.

The problem with this argument is that examination is performed by individual examiners whose effort is difficult for the patent office to monitor. While an examiner's identity is known

[37] The fact that application fees could be negative raises another issue: to collect the subsidy, firms might apply for things that are indeed novel and nonobvious, but have little or no value. One may therefore have to restrict fees to be nonnegative. Even with a nonnegativity constraint on fees, however, differentiation of fees would remain optimal.

to applicants, examiners are randomly assigned to applications within their field of expertise (firms cannot choose their examiner). Random assignment severely limits an examiner's ability to build a reputation for rigor.[38]

Tying examiners' payment to court rulings of validity

One reason why information is not entirely soft in practice is that examiners' decisions are subject to judicial review. This suggests a seemingly attractive way of improving examiners' incentives: basing their bonus payments on court rulings over applications they have handled. There are two problems that make this impractical, however. In many cases, court hearings occur years after the original patent examination. The examiner who was in charge of handling the application may no longer even work at the patent office.[39] Moreover, the courts' "patent friendliness" may evolve over time. Observers have suggested that this was the case in the United States after the creation of a centralized appeals court for patent disputes, the Court of Appeals for the Federal Circuit (CAFC).

An alternative option that would seem easier to implement in practice is to subject examiners to random peer review: applications would, with some probability, be sent to a second examiner, and both examiners' payment would depend on whether or not they concur.[40] Peer review needs to be carefully designed in order to avoid collusion.

The distribution of the signal

The distribution of the signal we have assumed is restrictive in two ways: there is a signal B but no signal R, and there are type-II but no type-I errors. We believe that this is a good description of patent examination. To prove that an invention is not novel, a patent examiner has to come up with prior art showing that the technology had already been described elsewhere or would have been obvious to someone of ordinary skill in the art. The prior art corresponds to the signal B in the model. By contrast, proving that an invention is novel would require showing that it has never been described anywhere else. Therefore, a signal R, especially if assumed hard information, would correspond to a description of the entire stock of knowledge in the world, which does not seem reasonable. Similarly, if an invention is truly new, it is impossible to find prior art that describes it. There is more leeway when it comes to determining obviousness, but we believe that this leeway is more accurately captured by the concept of soft information.

[38] Of course, this raises the question of why the assignment is random. Probably, issues of collusion also enter the equation here.

[39] This is especially true at the USPTO, which has a much higher turnover than the EPO. The average USPTO examiner stays for only three years (van Pottelsberghe and François, forthcoming).

[40] Mark Schankerman suggested such a procedure in a discussion with the author.

This being said, in the Appendix we consider a more general signal distribution allowing for both a signal R and type-I errors (i.e., rejecting a valid application). Not surprisingly, a signal R alleviates the commitment problem. Assuming that the signal is informative about effort, in the hard-information case a transfer for coming up with signal R provides incentives even if in equilibrium $p = 1$. Moreover, in the soft-information case, the planner can pay for a signal that is positively related to effort without jeopardizing truthful revelation. Nevertheless, the full-commitment outcome cannot be achieved: when p tends to 1, posterior beliefs for all signals converge, so incentive compatibility requires equal transfers for all signals. But equal transfers combined with equal posterior beliefs eliminate (extrinsic and intrinsic) incentives for effort provision. The problem caused by soft information is more fundamental than the restrictions we have imposed on the signal distribution.

6 Conclusion

We have presented a model of patent examination where applications are of uncertain validity and an agency is charged with determining patentability. Firms differ in creativity and self-select between doing genuine R&D, submitting bogus applications, and inactivity. Their behavior is determined by the application fee and the agency's examination effort. We have shown that the optimal full-commitment policy entails complete deterrence of bogus applications. The equilibrium of the examination game depends critically on whether the signal obtained by the patent office is hard or soft information. With hard information, an outcome that is arbitrarily close to the optimum can be achieved. With soft information, however, the planner must trade off innovation against deterrence, and the full-commitment outcome is unattainable. A concern for social welfare on the part of the agency is indispensable if it is to be induced to exert effort and truthfully reveal its signal.

We have also shown that the main results of the model are robust to changes in a number of assumptions. We now briefly discuss some limitations. First, soft information is neither necessary nor sufficient to have an incentive scheme that rewards the agency for grants. When the ex post welfare effects of patents are positive, or the agency cares about making (legally) correct decisions, it may be possible to pay it for rejections and still obtain truthful revelation. When the agency can produce a positive signal (i.e., one that is evidence of novelty), it can be optimal to pay it for grants under both hard and soft information. Second, for the full-deterrence result it is important that firms do not have wealth or liquidity constraints, so that an application fee that is large enough to deter bogus applications is affordable to those doing research. Finally, the results implicitly rely on the assumption that secrecy is ineffective as a protection mechanism, and equally so for all firms. In practice, the effectiveness of patents and secrecy is likely to vary across industries (Cohen *et al.*, 2000). If secrecy is a socially less desirable way of appropriating returns from innovation, the planner may want to encourage patenting. This places constraints on the use of the application fee as a deterrent.

While the results of the model should be taken with a grain of salt, it highlights the importance of soft information and produces insights that are policy relevant. To some extent, the nature of the information produced by the patent office can be affected by legislation. Vagueness in the definition of patentability criteria tends to move the signal towards soft information. The model suggests that efforts by legislators to reduce vagueness will have beneficial effects on the quality of patents by improving the incentives that can be designed for patent examiners.

The model also underscores the importance of selecting examiners who care about their job. How this can be achieved, however, is beyond the scope of this paper. In a recent contribution, Prendergast (2007) studies the sorting of workers with heterogeneous motivations into bureaucracy. He shows that bureaucrats should be biased, in the sense of not sharing the principal's preferences. Self-selection, however, does not generally lead to a workforce with the desired bias. A number of questions about the sorting of workers into bureaucracy remain unanswered and merit further research.

APPENDIX: A different signal distribution

Suppose that the distribution of the signal is as shown in table 1, where $y_{d\sigma}(e)$ is the probability of the patent office receiving signal σ conditional on effort e and on the applicant's activity d.

	$\sigma = R$	$\sigma = B$	$\sigma = \varnothing$
$d = R$	$y_{RR}(e)$	$y_{RB}(e)$	$1 - y_{RR}(e) - y_{RB}(e)$
$d = B$	$y_{BR}(e)$	$y_{BB}(e)$	$1 - y_{BR}(e) - y_{BB}(e)$

Table 1: Distribution of the signal conditional on the applicant's activity

The difference with the distribution in the text is that there is a second signal R and that the patent office can make both type-I and type-II errors. The basic model is a special case of this with $y_{RR}(e) = y_{RB}(e) = y_{BR}(e) = 0$ and $y_{BB}(e) = e$.

Assumption 6.1. *The probabilities $y_{d\sigma}(e)$ satisfy*

(i) $\dfrac{y_{RR}(e)}{y_{BR}(e)} \geq \dfrac{1 - y_{RR}(e) - y_{RB}(e)}{1 - y_{BR}(e) - y_{BB}(e)} \geq \dfrac{y_{RB}(e)}{y_{BB}(e)}$ *for all e;*

(ii) $y'_{RB} \leq 0$ *and* $y'_{BR} \leq 0$*;*

(iii) $y'_{RR} + y'_{RB} \geq 0$ *and* $y'_{BR} + y'_{BB} \geq 0$*.*

Assumption 6.1 says that regardless of the level of effort the distribution of σ satisfies the MLRP with respect to the applicant's activity (property (i)). Since MLRP implies first-order stochastic dominance, it follows that $y_{BB}(e) \geq y_{RB}(e)$ and $y_{RR}(e) \geq y_{BR}(e)$. Properties (ii) and (iii), which imply $y'_{RR} \geq 0$ and $y'_{BB} \geq 0$, mean that effort increases the probability of obtaining a correct signal and decreases both the probability of an incorrect signal and the probability of no signal.

Let $\hat{p}_\sigma(e)$ denote the posterior belief that the applicant has done genuine research given signal σ and effort e. MLRP implies that posterior beliefs are ordered as follows: $\hat{p}_R \geq \hat{p}_\varnothing \geq \hat{p}_B$. We prove here the first of these inequalities; the second can be proved in an analogous manner. By Bayes' rule, we have

$$\hat{p}_R = \frac{p y_{RR}}{p y_{RR} + (1-p) y_{BR}}$$
$$\hat{p}_\varnothing = \frac{p(1 - y_{RR} - y_{RB})}{p(1 - y_{RR} - y_{RB}) + (1-p)(1 - y_{BR} - y_{BB})},$$

After computation, we obtain that $\hat{p}_R \geq \hat{p}_\varnothing$ is equivalent to

$$y_{RR}(1 - y_{BB}) \geq y_{BR}(1 - y_{RB}),$$

which is true because, from the definition of MLRP, $y_{RR}(1-y_{BR}-y_{BB}) \geq y_{BR}(1-y_{RR}-y_{RB})$.

Define x_σ as the probability that an applicant is granted a patent if the patent office receives signal σ. Instead of going through the full-fledged optimization problem to determine grant probabilities, here we take the following, simpler approach: we impose only a mild consistency condition, namely $1 \geq x_R \geq x_\varnothing \geq x_B \geq 0$, and see whether the main results of the paper are robust.[41] For notational convenience, let $\mathbf{x} = (x_R, x_B, x_\varnothing)$ be the vector of grant probabilities and $\mathbf{y}_d(e) = (y_{dR}(e), y_{dB}(e), 1 - y_{dR}(e) - y_{dB}(e))$ be the vector of signal probabilities given activity d. A firm of type θ doing research obtains an expected payoff denoted $\Pi_R(\theta, e)$ while a firm submitting a bogus application obtains an expected payoff denoted $\Pi_B(\theta, e)$ (gross of application fees). Using the model of section 5.3, we have

$$\begin{aligned} \Pi_R(\theta, e) &= \theta[x_R y_{RR}(e) + x_B y_{RB}(e) + x_\varnothing(1 - y_{RR}(e) - y_{RB}(e))] - \psi(\theta) \\ \Pi_B(\theta, e) &= b\theta[x_R y_{BR}(e) + x_B y_{BB}(e) + x_\varnothing(1 - y_{BR}(e) - y_{BB}(e))], \end{aligned}$$

or, in matrix notation (letting superscript T denote the transpose),

$$\begin{aligned} \Pi_R(\theta, e) &= \theta \mathbf{x} \mathbf{y}_R^T(e) - \psi(\theta) \\ \Pi_B(\theta, e) &= b\theta \mathbf{x} \mathbf{y}_B^T(e), \end{aligned}$$

The innovation threshold $\hat{\theta}$ is now defined by $\Pi_R(\hat{\theta}, e) = \Pi_B(\hat{\theta}, e)$. To establish whether, even in this more general setup, greater examination effort by the patent office makes research relatively more attractive than imposture, we have to compute $\frac{\partial}{\partial e}[\Pi_R(\hat{\theta}, e) - \Pi_B(\hat{\theta}, e)]$, yielding

$$\frac{\partial}{\partial e}\left[\hat{\theta}\mathbf{x}\big(\mathbf{y}_R^T(e) - b\mathbf{y}_B^T(e)\big) - \psi(\hat{\theta})\right] = \hat{\theta}\left[(x_R - x_\varnothing)(y'_{RR} - by'_{BR}) + (x_B - x_\varnothing)(y'_{RB} - by'_{BB})\right]. \quad (36)$$

By properties (ii) and (iii) of Assumption 6.1 and under the condition that $x_R \geq x_\varnothing \geq x_B$, the expression in (36) is nonnegative, so e indeed (weakly) increases the relative profitability of research.

While it ensures that patent examination plays a role in encouraging R&D and avoiding rent-seeking, this does not guarantee the single-crossing condition that is needed for the full-deterrence result of Proposition 1. To have single crossing, we need $\frac{\partial}{\partial \hat{\theta}}[\Pi_R(\hat{\theta}) - \Pi_B(\hat{\theta})] \geq 0$ as well. This requires

$$\mathbf{x}\big(\mathbf{y}_R^T(e) - b\mathbf{y}_B^T(e)\big) \geq \psi'(\hat{\theta}). \quad (37)$$

Since by assumption ψ' is increasing, it is sufficient that this condition be satisfied as $\theta \to \infty$, but the restrictions on ψ' which are needed for the condition to hold are stronger than those given in Assumption 5.

[41] Given the ordering of posterior beliefs that MLRP implies, such a policy intuitively makes sense. Moreover, if instead of treating them as exogenous, we let the planner choose grant probabilities, this would give the planner an additional set of instruments and distort the comparison with the basic model.

We now turn to the equilibrium of the examination game, starting with the hard-information case. Suppose again that each examiner handles a single application so that transfers can condition only on the signal the examiner produces. Let $\mathbf{t} = (t_R, t_B, t_\varnothing)$ denote the vector of transfers associated to the signals. An additional issue that arises with the new signal distribution is that an examiner may want to hide a signal $\sigma = R$ and pretend not to have found any signal. To determine effort incentives, suppose the examiner anticipates that he will truthfully reveal his signal. The examiner chooses effort to maximize

$$p\mathbf{y}_R(e)[(1-\alpha)\mathbf{t}^T - \alpha D\mathbf{x}^T] + (1-p)\mathbf{y}_B(e)[(1-\alpha)\mathbf{t}^T - \alpha L\mathbf{x}^T] - \gamma(e),$$

leading to the first-order condition

$$p\frac{\mathrm{d}\mathbf{y}_R(e)}{\mathrm{d}e}[(1-\alpha)\mathbf{t}^T - \alpha D\mathbf{x}^T] + (1-p)\frac{\mathrm{d}\mathbf{y}_B(e)}{\mathrm{d}e}[(1-\alpha)\mathbf{t}^T - \alpha L\mathbf{x}^T] = \gamma'(e)$$

which can be rewritten as

$$\begin{aligned}(1-\alpha)\big[t_R[py'_{RR} + (1-p)y'_{BR}] + t_B[py'_{RB} + (1-p)y'_{BB}] - t_\varnothing[p(y'_{RR} + y'_{RB}) + \\ + (1-p)(y'_{BR} + y'_{BB})]\big] - \alpha\big[x_R[pDy'_{RR} + (1-p)Ly'_{BR}] + x_B[pDy'_{RB} + (1-p)Ly'_{BB}] - \\ - x_\varnothing[pD(y'_{RR} + y'_{RB}) + (1-p)L(y'_{BR} + y'_{BB})]\big] = \gamma'(e). \qquad (38)\end{aligned}$$

Equation (38) shows that under Assumption 6.1 effort reacts negatively to $t_\varnothing$, while its reaction to t_R and t_B is ambiguous in general. For large p, effort is increasing in t_R and decreasing in t_B. For small p, effort is increasing in t_B and decreasing in t_R. For intermediate values of p the direction is ambiguous. It is always increasing in at least one of the transfers (t_R and t_B), however. Suppose otherwise, i.e. that effort is decreasing in both t_R and t_B, or

$$py'_{RR} + (1-p)y'_{BR} < 0 \qquad (39)$$

$$py'_{RB} + (1-p)y'_{BB} < 0. \qquad (40)$$

We know from Assumption 6.1 that

$$p(y'_{RR} + y'_{RB}) \geq 0. \qquad (41)$$

Subtracting (41) from (39), we have

$$(1-p)y'_{BR} - py'_{RB} < 0. \qquad (42)$$

But adding up (42) and (40), we get

$$(1-p)[y'_{BR} + y'_{BB}] < 0,$$

a contradiction with Assumption 6.1. Thus, either (39) or (40) must be positive.

What this analysis tells us for the hard-information case is that

- the planner should not reward examiners for not coming up with a signal, so it is optimal to set $t_\varnothing = 0$;

- the transfers t_R and t_B should depend on p. The volume of applications can be used to measure p since in the model it is a perfect indicator of the share of good applicants;

- it is easier to attain the full-commitment outcome than in the basic model: $p = 1$ is now possible as an equilibrium;

- at the optimum, p will be close to or even equal to 1 and thus the patent office will be rewarded only for finding signal $\sigma = R$: $t_R > t_B = 0$.

The transfer t_R needs to respect an additional constraint which is due to the fact that it is always possible for the examiner to hide a signal and claim not to have found anything.[42] Defining $\Delta_\sigma(e) \equiv \hat{p}_\sigma(e)D + (1 - \hat{p}_\sigma(e))L$, it must be the case that

$$(1-\alpha)t_R - \alpha x_R \Delta_R(e) \geq (1-\alpha)t_\varnothing - \alpha x_\varnothing \Delta_R(e)$$

which, using the fact that $t_\varnothing = 0$, can be written as

$$(1-\alpha)t_R \geq \alpha(x_R - x_\varnothing)\Delta_R(e).$$

This constraint places a lower bound on t_R whenever $x_R > x_\varnothing$, but does not affect the qualitative results.

With soft information, there is now a larger number of incentive-compatibility constraints than in the basic model. The constraints for type $\sigma = R$ are

$$(1-\alpha)t_R - \alpha x_R \Delta_R(e) \geq (1-\alpha)t_\varnothing - \alpha x_\varnothing \Delta_R(e) \quad (43)$$

$$(1-\alpha)t_R - \alpha x_R \Delta_R(e) \geq (1-\alpha)t_B - \alpha x_B \Delta_R(e). \quad (44)$$

The constraints for type $\sigma = \varnothing$ are

$$(1-\alpha)t_\varnothing - \alpha x_\varnothing \Delta_\varnothing(e) \geq (1-\alpha)t_R - \alpha x_R \Delta_\varnothing(e) \quad (45)$$

$$(1-\alpha)t_\varnothing - \alpha x_\varnothing \Delta_\varnothing(e) \geq (1-\alpha)t_B - \alpha x_B \Delta_\varnothing(e). \quad (46)$$

The constraints for type $\sigma = B$ are

$$(1-\alpha)t_B - \alpha x_B \Delta_B(e) \geq (1-\alpha)t_R - \alpha x_R \Delta_B(e) \quad (47)$$

$$(1-\alpha)t_B - \alpha x_B \Delta_B(e) \geq (1-\alpha)t_\varnothing - \alpha x_\varnothing \Delta_B(e). \quad (48)$$

[42] The same is not true for t_B since $x_\varnothing \geq x_B$ and $t_\varnothing = 0$, so the examiner never prefers hiding a signal B.

From (44) and $x_R \geq x_\varnothing$, it follows that $t_R \geq t_\varnothing$. From (46) and $x_\varnothing \geq x_B$, it follows that $t_\varnothing \geq t_B$. Combining the two, we also have $t_R \geq t_B$. Moreover, as in the case of the basic model, whenever $x_\varnothing > x_B$ we must have $t_\varnothing > 0$, which by (38) reduces the incentives to provide effort. Unlike in the case of the basic model, effort responds positively to t_R (at least if p is sufficiently large). This tends to alleviate the incentive problem that soft information creates: the planner can now reward the agency for a signal that is positively related to effort. However, there is still an upper bound on the feasible incentive power. Combining (45) and (47), truthful revelation requires

$$t_R \leq \min\{t_\varnothing + \frac{\alpha}{1-\alpha}(x_R - x_\varnothing)\Delta_\varnothing(e), t_B + \frac{\alpha}{1-\alpha}(x_R - x_B)\Delta_B(e)\}.$$

Moreover, there is a fundamental problem that makes it impossible to achieve full deterrence when information is soft: as p approaches 1, $\Delta_\sigma \to D$ for all σ. Because posterior beliefs converge as the prior tends to 1, the expected social loss is the same whatever the signal. This means that the incentive-compatibility constraints cannot simultaneously be satisfied unless $t_R = t_\varnothing = t_B$. But then, there are no incentives to provide effort. Therefore, the result that soft information prevents the planner from attaining the full-commitment outcome remains valid.

References

Armstrong, Mark and Sappington, David E.M. (2007): Recent Developments in the Theory of Regulation. In: Mark Armstrong and Rob Porter (eds.), *Handbook of Industrial Organization*, Vol. 3, pp. 1557–1700. North-Holland: Elsevier.

Becker, Gary S. (1968): Crime and Punishment: An Economic Approach. *Journal of Political Economy*, Vol. 76(2), pp. 169–217.

Bessen, James and Maskin, Eric (forthcoming): Sequential Innovation, Patents, and Imitation. *RAND Journal of Economics.*

Bessen, James and Meurer, Michael J. (2008): *Patent Failure. How Judges, Bureaucrats, and Lawyers Put Innovators at Risk.* Princeton, NJ: Princeton University Press.

Caillaud, Bernard and Duchêne, Anne (2005): Patent Office in Innovation Policy: Nobody's Perfect. Working Paper, Paris School of Economics.

Chiou, Jing-Yuan (2006): The Design of Post-Grant Patent Challenges. Working Paper, Université de Toulouse 1.

Chiou, Jing-Yuan (2008): The Patent Quality Control Process: Can We Afford An (Rationally) Ignorant Patent Office? Working Paper, IMT Lucca.

Cohen, Wesley M., Nelson, Richard R., and Walsh, John P. (2000): Protecting Their Intellectual Assets: Appropriability Conditions and Why U.S. Manufacturing Firms Patent (or Not). NBER Working Paper no. 7552.

Cornelli, Francesca and Schankerman, Mark (1999): Patent Renewals and R&D Incentives. *RAND Journal of Economics*, Vol. 30, pp. 197–213.

Denicolò, Vincenzo and Franzoni, Luigi Alberto (2004a): The Contract Theory of Patents. *International Review of Law and Economics*, Vol. 23, pp. 365–380.

Denicolò, Vincenzo and Franzoni, Luigi Alberto (2004b): Patents, Secrets, and the First-Inventor Defense. *Journal of Economics and Management Strategy*, Vol. 13(3), pp. 517–538.

Farrell, Joseph and Shapiro, Carl (forthcoming): How Strong Are Weak Patents? *American Economic Review*.

Ford, George S., Koutsky, Thomas M., and Spiwak, Lawrence J. (2007): Quantifying the Cost of Substandard Patents: Some Preliminary Evidence. Phoenix Center Policy Paper no. 30.

Friebel, Guido, Koch, Alexander K., Prady, Delphine, and Seabright, Paul (2006): Objectives and Incentives at the European Patent Office. IDEI Report.

Green, Jerry R. and Scotchmer, Suzanne (1995): On the Division of Profit In Sequential Innovation. *RAND Journal of Economics*, Vol. 26(1), pp. 20–33.

Hopenhayn, Hugo, Llobet, Gerard, and Mitchell, Matthew (2006): Rewarding Sequential Innovators: Prizes, Patents, and Buyouts. *Journal of Political Economy*, Vol. 114(6), pp. 1041–1068.

Hopenhayn, Hugo A. and Mitchell, Matthew F. (2001): Innovation Variety and Patent Breadth. *RAND Journal of Economics*, Vol. 32(1), pp. 152–166.

Iossa, Elisabetta and Legros, Patrick (2004): Auditing and Property Rights. *RAND Journal of Economics*, Vol. 35(2), pp. 356–372.

Jaffe, Adam B. and Lerner, Josh (2004): *Innovation and Its Discontents: How Our Broken Patent System is Endangering Innovation and Progress, and What to Do About It.* Princeton, NJ: Princeton University Press.

Kim, Son Ku (1997): Limited Liability and Bonus Contracts. *Journal of Economics and Management Strategy*, Vol. 6(4), pp. 899–913.

Langinier, Corinne and Marcoul, Philippe (2003): Patents, Search of Prior Art and Revelation of Information. Working Paper, Iowa State University.

Lemley, Mark A. (2001): Rational Ignorance at the Patent Office. *Northwestern University Law Review*, Vol. 95(4), pp. 1497–1532.

Melumad, Nahum D. and Mookherjee, Dilip (1989): Delegation as Commitment: the Case of Income Tax Audits. *RAND Journal of Economics*, Vol. 20(2), pp. 139–163.

Merges, Robert P. (1999): As Many As Six Impossible Patents Before Breakfast: Property Rights for Business Concepts and Patent System Reform. *Berkeley Technology Law Journal*, Vol. 14(2), pp. 577–616.

Perry, James L. and Wise, Lois R. (1990): The Motivational Bases of Public Service. *Public Administration Review*, Vol. 50(3), pp. 367–73.

Prady, Delphine (2008): 'Is That a Gun In Your Pocket?' Overadvertising and Evaluation. Working Paper.

Prendergast, Canice (2007): The Motivation and Bias of Bureaucrats. *American Economic Review*, Vol. 97(1), pp. 180–196.

Scotchmer, Suzanne (1999): On the Optimality of the Patent Renewal System. *RAND Journal of Economics*, Vol. 30(2), pp. 181–196.

Shin, Hyun Song (1998): Adversarial and Inquisitorial Procedures in Arbitration. *RAND Journal of Economics*, Vol. 29(2), pp. 378–405.

Strausz, Roland (1997): Delegation of Monitoring in a Principal-Agent Relationship. *Review of Economic Studies*, Vol. 64, pp. 337–357.

van Pottelsberghe, Bruno and François, Didier (forthcoming): The Cost Factor in Patent Systems. *Journal of Industry, Competition and Trade.*