Inventors and Impostors: An Economic Analysis of Patent Examination*

Florian Schuett†

August 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 invented something which is not truly novel). The government has to delegate the verification of novelty to an agency which must exert costly effort in order to obtain a signal of validity. Firms self-select depending on 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 depends critically 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, consistent with results from the literature on optimal audits. We argue, however, that in the case of patent examination the signal is unlikely to be verifiable. With soft information, the government is constrained in the incentives it can offer the agency because transfers need to be designed to insure truthfulness. Therefore, there is a tradeoff between the sustainable examination effort and deterrence of bogus applications.

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

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

---

* I thank Vincenzo Denicolò, Guido Friebel, Mark Schankerman, Paul Seabright, Jean Tirole, participants at the EVPAT summer school in Bologna and the "Knowledge for Growth" conference in Toulouse, as well as seminar participants at Toulouse School of Economics for helpful comments and suggestions. All errors are mine.

† Toulouse School of Economics (GREMAQ), Manufacture des Tabacs, Aile Jean-Jacques Laffont, Office MF004, 21 allée de Brienne, 31000 Toulouse (France), Email: `florian.schuett@univ-tlse1.fr`

1 Introduction

To be patentable, an invention must satisfy the standards of novelty and non-obviousness. Many issued patents, however, do not satisfy these standards. Such so-called "bad patents" arguably cause substantial social losses: they lead to costly litigation, cause deadweight loss, and create disincentives to future innovators. Ford *et al.* (2007) estimate the total cost of bad patents to the US economy at $25.5 billion annually.[1] As several prominent firms such as eBay and R.I.M., the maker of Blackberry mobile devices, have been drawn into infringement lawsuits by holders of dubious patents,[2] the entire process of patent examination has come under increased public scrutiny.[3] Why does the patent office grant such a large number of bad patents? Lemley (2001) has argued 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. Screening out all bad applicants would be extremely costly without providing much offsetting benefit. This "rational ignorance" argument is uncontroversial in principle. While some balance between benefits and costs of more thorough examination must be struck, many observers feel, though, that the U.S. Patent and Trademark Office (USPTO) has not got the balance right and is being too ignorant (Jaffe and Lerner, 2004; Farrell and Shapiro, forthcoming).

An alternative explanation, starting from the idea that patent examination resembles an inspection game, attributes the patent office's apparent laxity to commitment problems. The office would like to convince applicants that it will be extremely rigorous in screening their applications. But if applicants believed this, none of them would submit invalid applications, and the patent office would no longer have an incentive to perform a rigorous screen. In equilibrium, so the argument goes, there will be some bad applications, and the patent office's examination effort will not suffice to screen out all of them. Results from the auditing literature, however, cast doubt on the plausibility of this explanation. Melumad and Mookherjee (1989) have shown that delegation of auditing decisions to a salary-maximizing manager can achieve the full-commitment outcome. Do these insights carry over to the case of patent examination? That is, can a patent examiner who responds to monetary incentives be induced to choose the optimal level of screening in equilibrium, even though it is ex post inefficient?

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

[2] RIM, for example, was sued by a company called NTP and settled out of court for a reported $612.5 million, even though the US Patent and Trademark Office (USPTO) rejected 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.

[3] See, e.g., the books by Jaffe and Lerner (2004) and Bessen and Meurer (2008). The US Congress is currently deliberating on an extensive patent reform bill.

In this paper, we show that delegation of patent examination to an agency can indeed solve the commitment problem, provided that the agency's signal about patentability is hard information. This confirms the results obtained by Melumad and Mookherjee (1989) and extends them to a moral-hazard framework. We go on to argue, though, that the agency's signal is unlikely to be hard information. When the signal is soft information, it is no longer possible to achieve the full-commitment outcome. In equilibrium, patent examination is less rigorous than it would be under the optimal policy.

In the model presented in section 2, the government delegates patent examination to an agency motivated by both monetary transfers and a concern for social welfare. The agency must expend effort in order 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 to get away with imposture, and therefore increases the attractiveness of true research. This setup leads to self-selection of firms, with high-creativity firms doing genuine research, and low-creativity firms submitting bogus applications. The rigor of patent examination determines the share of firms doing research. Our formulation thus 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, the government induces a level of effort that equalizes 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 ever issue. In the absence of commitment, this feature of the optimal policy clearly 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] 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 under the simplifying assumption that each examiner handles a single patent application. 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 rejecting applications if it comes up with invalidating prior art. The application fee is adjusted to achieve the optimal amount of innovation.

With soft information, the incentive scheme needs to ensure truthfulness on the part of the agency: since the agency will be tempted to reject even valid applications, incentives must be designed to make sure that rejections only occur when the agency has actually found invalidating prior art. This requires that the agency be paid for granting patents. The government's only remaining effective instrument is the application fee, which implies that 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 genuine research.

In section 5, we repeat the previous analysis assuming that each examiner handles multiple applications. The results from the hard-information case are strengthened: the government can now approximate the full-commitment outcome without leaving any rent to the examiner. This is achieved by means of a simple bonus scheme. The results from the soft-information case also prove to be robust. Even with multiple applications per examiner, the constraints imposed by the necessity to insure truthful revelation force the government to reward the examiner for granting patents. This dilutes his incentives to provide effort.

Section 6 discusses whether hard or soft information is a better description of reality, compares the predictions of the model with empirical observations, and elaborates on possible limitations. 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. The incentive scheme that ensures truthfulness in the soft-information case is 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 (Jaffe and Lerner, 2004). In addition, 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 (Dixit, 2002), and of examiners at the European Patent Office (EPO) in particular (Friebel *et al.*, 2006).

Related literature

The literature on incentives to innovate has used a mechanism-design approach to analyze the optimal structure of innovation policy (see, e.g., Cornelli and Schankerman, 1999; Scotchmer, 1999; Hopenhayn and Mitchell, 2001; Hopenhayn *et al.*, 2006). In doing so, it implicitly assumes that, for policy implementation, the mechanism designer can rely on a benevolent agency with commitment power. However, just as a regulator may be tempted to engage in opportunistic behavior, the patent office may also have limited commitment.[5] While the literature cited above assumes that the inventor has private information on the value of his innovation, it does not examine the patentability dimension, in spite of the attention it has received in the public discussion. If the government needs to give discretion over patentability decisions to an agency, and this agency has to exert effort to learn about novelty of applications, then opportunism with respect to the thoroughness of examination arises naturally.

Two other papers have investigated patent examination. 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 incentives to search if it can identify valid applications beforehand. While Langinier and Marcoul (2003) assume that the inventor is initially uninformed about the patentability of his innovation, Caillaud and Duchêne (2005) examine the case where the inventor knows whether or not he has a patentable invention. In their model, valid inventions stem from successful R&D projects and invalid ones from failed projects. Unlike in our model, inventors are ex ante identical and do not choose whether to do genuine research or not. The focus in Caillaud and Duchêne (2005) is 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 of these papers assume that the patent office is fully benevolent but lacks com-

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

mitment power. Thus, they do not take into account the result obtained by Melumad and Mookherjee (1989) in the context of auditing, according to which a benevolent planner can overcome the commitment problem by delegating patent examination and setting up an appropriate incentive scheme. 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.

Our paper is also related to Prendergast (2003) who is concerned with the efficiency of bureaucracies, and to Iossa and Legros (2004) who study auditing with soft information. Prendergast (2003) models a bureaucrat as someone controlling an allocation to a customer, and shows that bureaucracies are only used when customers cannot be trusted to choose the welfare-maximizing allocation. This mirrors the relationship between patent office and applicants in our model. Unlike in our model, however, in Prendergast (2003), customers can file complaints, affecting the bureaucrat's payoff, and bureaucratic rules are designed to avoid that bureaucrats "capitulate" to customers. Iossa and Legros (2004) show that under soft information, a necessary condition for the auditor to exert any effort is that he be given a stake in the project to be audited. Similarly, we show in the soft-information case that positive effort will only occur if the examiner is intrinsically motivated – that is, if he has a "stake" in the outcome, i.e., the resulting level of social welfare.

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 (under patent protection or not) 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, in either case, 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 precision of the signal depends on the agency's examination effort. If the application is for a genuine invention, there is no signal ($\sigma = \varnothing$). For bogus applications 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.[7] An important question is whether $\sigma = B$ is hard information, in the sense of being verifiable for 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 $\gamma(0) = \gamma'(0) = 0$).

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. Importantly, we assume that *both* matter to the agency.

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. 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, 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

[6] This can be seen as a reduced-form profit function resulting from a firm's investment choice; see footnote 8 below.

[7] This setup, where the agency can only err in one direction, can be justified if $\sigma = B$ is interpreted as the examiner coming up with "prior art" showing that the invention is not new or would have been obvious to someone of ordinary skill in the art.

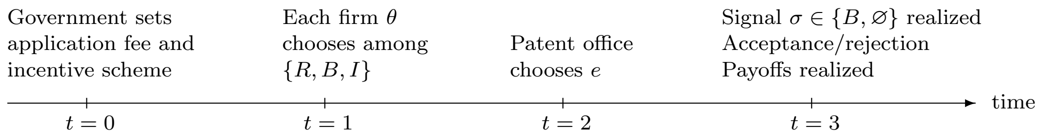

Figure 1: Timing of the game

first two, it is not clear whether patents are ex post welfare-enhancing or reducing. For the purposes of this paper, however, we assume $D \geq 0$, that is, patents do not increase welfare from an ex post point of view.

A patent on a bogus application causes a social loss of $L > 0$. As in the case of a true invention, this includes 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 tends to 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. We will therefore assume that $L > D$: the ex post loss from a valid patent is smaller than the loss from a bad patent. For simplicity, D and L are also assumed to be independent of θ. Appendix B develops a more structural model of the innovation process. The model accounts in a simple way for secrecy, sequential innovation, and court challenges. We derive conditions under which profit and welfare functions exhibit the assumed properties.

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. 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 make their decision.

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 $(1-e)\pi_B(\theta) - F$.[8] Thus, a firm prefers R&D to imposture if and only if

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

Assumption 2. *Profit functions satisfy*

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

(ii) $\pi'_R > \pi'_B > 0$*,*

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

For firms at the lower bound of the creativity distribution ($\theta = 0$), obtaining a patent on a bogus application is more profitable than producing a true invention. Profits from both activities increase with θ, perhaps because identifying valuable bogus applications requires some of the same qualities as identifying valuable research fields. Profits from research are more sensitive to creativity than those from bogus patents, though. Finally, first derivatives of the profit functions satisfy monotonicity conditions.[9] 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

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

such that firms with creativity higher than $\hat{\theta}$ do research, while 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.

[8] 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 application fee influence the optimal R&D investment, which would make the above analysis invalid. However, given the model setup, the level of investment, and thus π_R, is 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 his 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 $\pi_R(\theta) = \rho(z^*(\theta),\theta) - \psi(z^*(\theta))$.

[9] Conditions for these assumptions to hold are discussed in Appendix B.

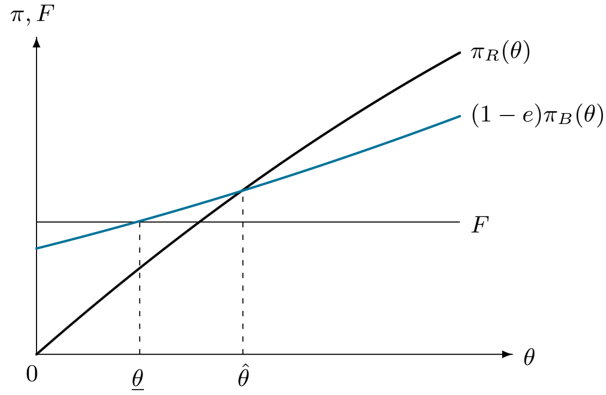

Figure 2: Self-selection of firms according to θ

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[10]

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

subject to (1), (2) and $\underline{\theta} \leq \hat{\theta}$. The last constraint is due to 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 2 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: See Appendix A.

Higher examination effort increases the attractiveness of genuine research relative to imposture. That is, e determines the incentives to do R&D.[11] The planner chooses e^o to equalize

[10] The objective function relies on costs of public funds being zero. Transfers to the agency and the amount of application fees raised do not enter, so we can ignore the agency's individual rationality constraint. In Appendix C, we consider the case where there is a positive shadow cost of public funds and derive a condition under which the main result of Proposition 1 (full deterrence) is robust.

[11] This result is robust to the assumption that genuine research never leads to unpatentable inventions.

marginal social gains from more innovation (the left-hand side of (4)) with the marginal cost of examination (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 is never any bogus application, and no bad patent is ever 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 to raising the fee up 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 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 the issuance of invalid patents and associated social losses. 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 may solve the principal's commitment problem (Melumad and Mookherjee, 1989; Strausz, 1997). In the following sections, 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 Incentive schemes with a single application per examiner

This section studies the simple case where there is one patent examiner for every applicant. This assumption, which reduces the set of incentive schemes available to the planner, is not without loss of generality but has the merit of conveying most of the intuition in a straightforward manner. In section 5, we study the more general case where each examiner treats many patent applications. 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 (the agency

Even if there were a chance of research leading to something that is not new, the level of examination effort would not decrease the incentives for R&D. What matters for the decision between research and imposture is the relative attractiveness of each of these activities. Increasing e still makes research relatively more attractive than imposture because the probability of having a valid invention is strictly greater when doing research.

cannot concoct false evidence) in section 4.1. In section 4.2 we turn to the case where the signal is soft, implying that the agency also has discretion over refusing valid applications.

4.1 Hard information

We solve the game backward, starting with the agency's effort choice. Suppose the agency believes that a proportion p of all applicants has patentable inventions.[12] By screening out a bad application, it avoids a social loss of L. The incentive scheme chosen by the planner can condition on two events: either the agency comes up with defeating prior art ($\sigma = B$), in which case it receives transfer t_r and the application is rejected, or it does not ($\sigma = \varnothing$), in which case it receives t_g and the patent is granted. The agency chooses e to maximize

$$p[(1-\alpha)t_g - \alpha\, D] + (1-p)[e(1-\alpha)t_r + (1-e)[(1-\alpha)t_g - \alpha\, L] - \gamma(e).$$

With probability p, it faces a valid applicant, so that it cannot find any grounds for rejection. The transfer it receives is t_g. 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_r 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

$$(1-p)\Big[\alpha\, L + (1-\alpha)(t_r - t_g)\Big] = \gamma'(e). \tag{5}$$

It is immediate from (5) 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, 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*

$$p^* = \frac{1 - G(h(e^*))}{1 - G(\ell(e^*, F))} \tag{6}$$

$$e^* = (\gamma')^{-1}\left((1-p^*)\big[\alpha\, L + (1-\alpha)(t_r - t_g)\big]\right). \tag{7}$$

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

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

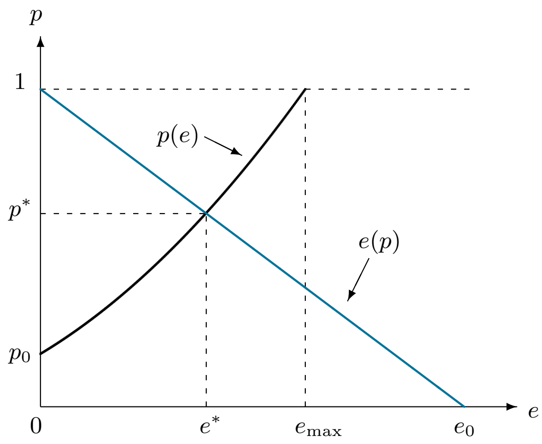

Figure 3: The equilibrium in the hard-information case

Proof: We prove existence and uniqueness of the stated equilibrium candidate. The agency's best-response function, $e(p)$, is obtained from (5). Firms' best response to examination effort e is

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

leading to a probability $p(e)$ as stated in (6), 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 $\alpha\, L + (1-\alpha)(t_r - t_g) > 0$ (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 $e_0 \equiv (\gamma')^{-1}\,(\alpha\, L + (1-\alpha)(t_r - t_g))$. 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 depicts the best-response functions for the examination game. The planner's choices of application fee and incentive scheme influence the equilibrium of the examination 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, $e_{\max}$. 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, $t_r - t_g$, affects the slope of the agency's best response function: the greater $t_r - t_g$, the larger (in absolute value) the slope. Summarizing the comparative statics,

- increasing $t_r - t_g$ 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 for *rejecting* applications. Since a signal that the application is not novel ($\sigma = B$) is assumed verifiable, the agency has to come up with actual evidence of invalidity to be able to refuse an application, meaning that it has to exert effort to obtain the reward.

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 the closer to $e_{\max}$ the larger is $t_r - t_g$. This means that, by setting $t_r - t_g$ 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 $t_r - t_g$ sufficiently large, the government can achieve an arbitrarily close approximation of the full-commitment outcome: for any ε there exists $t_r - t_g$ 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.

One potential caveat with this result is that the difference between transfers for acceptance and rejection may have to be quite large. Since there is a lower bound on transfers (limited liability), this may mean that a very high reward must be paid for rejecting an application. If public funds are costly or if there are budgetary restrictions, it may not be optimal (or feasible) to pay the agency a reward large enough to deter all bogus applications.[13] Moreover, the result that intrinsic motivation does not matter no longer holds, either: an intrinsically motivated agency saves on public funds.

It is in this context that the assumption of a single application per examiner is important. If each patent examiner handles more than one application, the planner can condition transfers

[13] This is true unless the planner can extract the rent from the agency through an appropriate upfront payment. The literature, however, generally interprets limited liability as ruling out negative first-period transfers.

on the total number of applications handled and on the total number of bad signals ($\sigma = B$) found. We examine this case in section 5 below.

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_r and rejects if the report is $\sigma = B$. She pays t_g and grants if the report is $\sigma = \varnothing$.[14]

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 over granting a patent, it must be the case that

$$(1-\alpha)t_r \geq (1-\alpha)t_g - \alpha\, L. \tag{8}$$

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

$$(1-\alpha)t_r \leq (1-\alpha)t_g - \alpha[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L] \tag{9}$$

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 rule this out we need

$$(1-\alpha)t_r \leq (1-\alpha)\big[[p+(1-p)(1-e^{**})]t_g+(1-p)e^{**}t_r\big]-\alpha[pD+(1-p)(1-e^{**})L]-\gamma(e^{**}). \tag{10}$$

Without making further assumptions on the functional form of γ, we cannot say much about how (10) relates to the other constraints. A sufficient condition for (10) to hold, however, is

$$(1-\alpha)t_r \leq (1-\alpha)t_g - \alpha[pD + (1-p)L], \tag{11}$$

[14] 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 incentives to provide effort. It can be shown that the second effect dominates the first.

as we show in the proof of Proposition 4. In what follows, we replace constraint (10) by (11).[15]

Since by Bayes' rule $\hat{p}(e) = p/[p + (1-p)(1-e)] > p$ for any $e > 0$ and, by assumption, $D < L$, (11) implies (9). Combining conditions (8) and (11), we obtain

$$pD + (1-p)L \leq \frac{1-\alpha}{\alpha}(t_g - t_r) \leq L. \tag{12}$$

To satisfy this condition, transfers must be such that $t_g \geq t_r$, i.e., the patent office is rewarded for *granting* patents. This is in stark contrast to the hard-information case examined in section 4.1 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

$$(1-p)[\alpha\, L + (1-\alpha)(t_r - t_g)] = \gamma'(e).$$

Since for any $\alpha > 0$, $t_r < t_g$ by (12), an incentive scheme that satisfies the condition for truthful grant/refusal decisions tends to reduce effort. Hence, the planner will choose $t_g - t_r$ at the smallest value consistent with (12), that is,

$$t_g - t_r = \frac{\alpha}{1-\alpha}[pD + (1-p)L]. \tag{13}$$

Thus, the planner's hands are tied as to the choice of transfers. She can no longer influence the slope of 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 a higher incidence of bogus applications and, ultimately, invalid patents.

Proposition 4. *With soft information, the planner is constrained to setting* (t_g, t_r) *as specified in* (13)*. 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 having no examination effort.

Proof: We start by showing that (11) is sufficient for (10) to hold. Rewrite (10) as

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

Since e^{**} satisfies (5) and γ is convex,

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

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

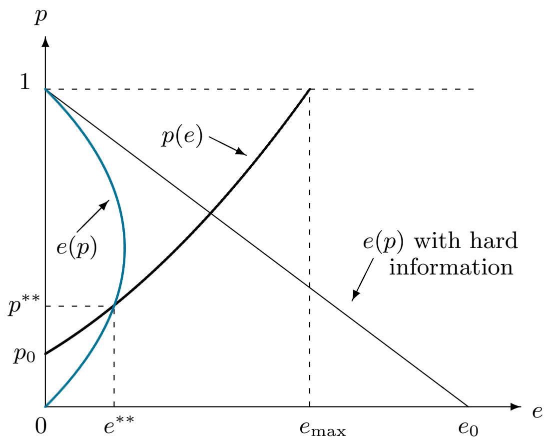

Figure 4: The equilibrium in the soft-information case

Thus, a sufficient condition for (10) is

$$(1-\alpha)[p+(1-p)(1-e^{**})](t_g-t_r)-\alpha[pD+(1-p)(1-e^{**})L] \geq (1-p)e^{**}[\alpha L+(1-\alpha)(t_r-t_g)]$$

which can be simplified to (11).

Plugging (13) into (5) 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 (6). Combining $e(p)$ and $p(e)$, one obtains the claimed result. Existence of equilibrium is guaranteed by $p(0) > 0$ which is true by Assumption 2 (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 $(1-p)p\alpha(L-D) < (1-p)[\alpha L+(1-\alpha)(t_r-t_g)]$ for any pair of transfers verifying $t_r \geq t_g$. ■

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_r and t_g are determined by (13), 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 upward or downward sloping part of the $e(p)$ curve. 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.

5 Incentive schemes with multiple applications per examiner

In this section, we relax the assumption that each examiner handles only a single patent application. 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.[16] 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 r; the number of grants is thus given by $a - r$. 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.

We will be interested in the two main results derived in section 4:

- In the hard-information case, we have shown that the planner can approximate the full-commitment outcome arbitrarily closely. But because of limited liability, this comes at

[16] 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.

the expense of a very large reward for rejections. The examiner obtains a rent which would be socially costly in the presence of a positive shadow cost of public funds. Can we achieve the efficient outcome without leaving a rent to the examiner? (Section 5.1.)

- In the soft-information case, 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? (Section 5.2.)

5.1 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 $\hat{r}$, and zero otherwise. That is,

$$t(r) = \begin{cases} 0 & \text{for } r < \hat{r} \\ T & \text{for } r \geq \hat{r} \end{cases} \tag{14}$$

Both $\hat{r}$ 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 $(\hat{r}, T)$ 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 very low target $\hat{r}$. If the number of bogus applications is small, the examiner will have to exert considerable effort to reach the 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 r 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 r is distributed on the support $[0, a]$ according to a continuous distribution $Q(r; e, a)$, with density $q(\cdot)$.[17]

Assumption 3. *The distribution of r has the following properties:*

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

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

(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 r 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 r. 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 r at a decreasing rate. Finally, property (iv), based on Kim (1997), makes sure that within the feasible set $[0, a]$ there exists a target $\hat{r}$ that permits implementation of a level of effort close to the full-commitment level through incentive scheme (14).

The examiner facing incentive scheme $t(r)$ from (14) solves the following problem:

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

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

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

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

Thanks to property (iii) in Assumption 3, 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 $(\hat{r}, T)$, implementing e^o_ε and leaving no rent to the examiner, must satisfy the following pair of equations for any $a > \bar{a}$:

$$[1 - Q(\hat{r}; e^o_\varepsilon, a)]T = a\gamma(e^o_\varepsilon) \tag{15}$$

$$-Q_e(\hat{r}; e^o_\varepsilon, a)T = a\gamma'(e^o_\varepsilon) \tag{16}$$

Equation (15) is the examiner's binding individual-rationality constraint, while (16) is the incentive-compatibility constraint.

Proposition 5 (Kim, 1997)**.** *Suppose Assumption 3 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 $t(r)$ from (14) such that $(\hat{r}, T)$ solve equations (15) and (16).*

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 (14) 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 $(\hat{r}, T)$ solving (15) and (16). From (15), $T = a\gamma(e^o_\varepsilon)/[1 - Q(\hat{r}; e^o_\varepsilon, a)]$. Plugging this into (16) and rearranging, we have

$$-\frac{Q_e(\hat{r}; e^o_\varepsilon, a)}{1 - Q(\hat{r}; e^o_\varepsilon, a)} = \frac{\gamma'(e^o_\varepsilon)}{\gamma(e^o_\varepsilon)}. \tag{17}$$

Under MLRP, the left-hand side is nondecreasing in $\hat{r}$ by Lemma 1 in Kim (1997). It is zero at its lower bound, $\hat{r} = 0$ 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 $r \to a$ by l'Hospital's rule. By property (iv) in Assumption 3, $q_e/q(a; e^o_\varepsilon, a) > \gamma'(e^o_\varepsilon)/\gamma(e^o_\varepsilon)$ when ε is small. Hence, an $\hat{r}$ solving (17) 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 3) and shown that, if any such contract exists, there exists a bonus scheme that implements the first-best allocation. In the proof of Proposition 5, the appropriate target $\hat{r}$ 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 $\hat{r}$ 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 5 thus strengthens the result of Proposition 3 according to which, with hard information, delegation can overcome the commitment problem. It demonstrates that the result does not hinge on the assumptions that public funds are costless and that each examiner handles only one application. In fact, with each examiner handling more than one application, the full-commitment outcome is easier to attain: unlike in the single-application case, it does not require leaving a rent to the examiner.[18]

5.2 Soft information

We now turn to the case of soft information. It should be clear that the incentive scheme in (14) 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 $\hat{r}$. 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 $\big(\tilde{r}, t(\tilde{r})\big)$, $\tilde{r} \in [0, a]$, 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 $t(\tilde{r})$ 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 r bad signals (r is his "type"). His utility from reporting $\tilde{r}$, gross of effort cost $a\gamma(e)$, is given by

$$U(\tilde{r}, r) = \begin{cases} (1-\alpha)t(\tilde{r}) - \alpha\,[\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L](a-\tilde{r}) & \text{for } \tilde{r} \geq r \\ (1-\alpha)t(\tilde{r}) - \alpha\,\big([\hat{p}(e^{**})D + (1-\hat{p}(e^{**}))L](a-r) + L(r-\tilde{r})\big) & \text{for } \tilde{r} < r \end{cases} \tag{18}$$

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 $\tilde{r}$, the examiner obtains a payoff from the

[18] 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.

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 $\tilde{r}$ times zero (rejected applicants do not cause any social loss) minus the expected social loss from the $a-\tilde{r}$ accepted applicants. When over-reporting ($\tilde{r} > r$), 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 $r-\tilde{r}$ accepted applications are invalid for sure, while for the remaining $a-r$, his posterior belief about validity is again $\hat{p}$.

Incentive compatibility (IC) requires

$$U(r,r) \geq U(r',r) \quad \forall (r,r') \in [0,a]^2. \tag{19}$$

In particular, (19) must hold for $r' = a$. 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 $\tilde{r} = a$.

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

$$E[U(r,r)|e^{**}] - a\gamma(e^{**}) \geq 0, \tag{20}$$

where the expectation is taken over r. We again resort to the use of a sufficient condition to replace (20), which is that it holds for $e = 0$,[19]

$$\begin{aligned} & E[U(r,r)|0] \geq 0 \\ \Leftrightarrow\ & t(0) \geq \frac{\alpha}{1-\alpha}[pD+(1-p)L]a. \end{aligned} \tag{21}$$

The following proposition characterizes incentive-compatible transfer schemes.

Proposition 6. *Suppose the signal is soft information and each examiner handles a applications. An optimal transfer scheme inducing truthful revelation must satisfy*

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

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

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

Proof: See Appendix A.

The optimal transfer scheme is again decreasing with r: 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 $t(r) = 0\ \forall r$ (or, more generally, any transfer that is independent of r). 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.[20] The full-commitment outcome cannot be attained, and the planner faces a tradeoff in choosing F: a lower F means more bogus applications, but also greater equilibrium effort and thus more innovation.

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.

6 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

[20] This is true qualitatively speaking. Equation (5) determining effort provision will of course have to be adapted to the multiple-application case.

evaluate whether a number of fictitious inventions satisfied different definitions of a "technical contribution" (Friebel *et al.*, 2006).[21] 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.

It is interesting to note that the soft-information model delivers results which seem 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 resulting incidence of bad patents;
- 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 (Jaffe and Lerner, 2004);
- the high degrees of intrinsic motivation displayed by EPO examiners, as documented in a survey by Friebel *et al.* (2006). The phenomenon is well known from other bureaucracies as well (Dixit, 2002).[22] It 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.[23]

Reality, of course, is likely to lie somewhere between hard and soft information. 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 trials 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 (especially at the USPTO which has much higher turnover than the EPO). Moreover, the courts' "patent friendliness" may evolve over time. Observers have suggested that this was the case in the

[21] 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.

[22] See Prendergast (2007) for a general model of bureaucracies incorporating intrinsically motivated bureaucrats.

[23] At most, intrinsic motivation relaxes the examiners' participation constraint.

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 seems 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.[24] Peer review needs to be carefully designed in order to avoid collusion.

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 to monitor for the patent office. While an examiner's identity is known 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.[25]

The nature of intrinsic motivation

A related point concerns the way we have defined intrinsic motivation. Examiners are assumed to care about social welfare, but since their decisions take place after R&D costs are sunk, they do not take into account the effects of their effort choice on incentives to innovate. By contrast, bureaucrats' intrinsic motivation is often interpreted as "taking pride in doing the job right." The job of a judge, for example, is to apply the law. 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.

A judge, however, is subject to much greater public scrutiny than a patent examiner, especially in high-profile cases. The quality of his decisions (measured in terms of correctly

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

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

applying the law) also plays a much larger role in shaping his career.[26] For these reasons, it is difficult to disentangle image and career concerns on the one hand, and intrinsic motivation on the other hand, as drivers of a judge's decision making. It seems safe to say that the jury is still out on what motivates bureaucrats: adhering to their job description, or serving the greater public good.

Unresolved issues and future research

In the model, we have assumed 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. A uniform application fee also rules out renewal fees, i.e., differentiating fees according to the desired length of patent protection. In future research, we would like to understand how these two instruments (fines and renewal fees) can be used to deter bogus applicants without jeopardizing genuine research.

The absence of fines in practice may be related to the possibility that even firms doing genuine research may sometimes inadvertently re-invent old ideas. Penalizing them too heavily would reduce incentives to engage in research. Alternatively, collusion between examiners and applicants, or their competitors, may be a part of the puzzle: fines tend to increase the bargaining power of examiners. A more detailed exploration of these issues, however, must be left to future research.

7 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 three activities: 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 policy with observable effort entails full deterrence of bogus applications. When the agency lacks commitment power and its effort is unobservable, the outcome depends critically on whether or not 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 examination effort against invalid patents, and the full-commitment outcome is unattainable.

[26] Career concerns play a particularly small role at the USPTO, where turnover is high and the average examiner stays for all of three years (van Pottelsberghe and François, forthcoming).

Appendix A Omitted proofs

Proof of Proposition 1:

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

$$\frac{\partial \ell}{\partial F}\left[g(\underline{\theta})[(1-e)L+\gamma(e)]-\mu\right]=0. \tag{22}$$

Since $\partial\ell/\partial F > 0$, $\mu > 0$, so indeed $\underline{\theta} = \hat{\theta}$. We obtain (4) by differentiating (3) with respect to e, substituting for μ from (22) 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 2. This completes the proof. ■

Proof of Proposition 5:

Notice first that the utility function in (18) 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

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

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 (20), and thus never greater than $\frac{\alpha}{1-\alpha}[pD + (1-p)L]a$ by (21)), and the fact that, by property (ii), $t(r)$ is strictly decreasing in r. ■

Appendix B A more structural model

In the main text 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, infinitely-lived consumer with additively separable preferences who consumes one unit of each available good.[27] 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 $1-b$, in which case, again, any firm can produce it.[28] 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.

[27] 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.

[28] 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 doing genuine research obtains a profit (gross of patent application fees) $\pi_R(\theta) = \theta - \psi(\theta)$, while a firm filing a bogus application obtains an expected profit of $\pi_B(\theta) = b\theta$. We now give conditions on $\psi(\cdot)$ which make sure that π_R and π_B satisfy Assumption 2.

Assumption B.1. *The cost of research satisfies*

(i) $\psi(0) > 0$

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

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

Since $\pi_B(0) = 0$ and $\pi_R(0) = -\psi(0)$, having $\pi_B(0) > \pi_R(0)$ requires $\psi(0) > 0$. Convexity of $\psi(\cdot)$ is equivalent to concavity of $\pi_R(\cdot)$. Finally, 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.

Having established conditions under which profit functions fit the assumptions of the main text, we now turn to welfare. In the basic 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 $\psi(v)$.[29] 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 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 the 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

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

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

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

If the same innovator is refused a patent, he still has the possibility to protect his innovation through secrecy.[30] The corresponding welfare effect is

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

[29] 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).

[30] 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.

since the second innovation can only occur if the first one leaks out. The difference between the two expressions corresponds to D as defined in the main text,

$$D = S(\theta) - W(\theta) = \Phi - sV. \tag{23}$$

Equation (23) highlights the tradeoffs associated with patents after innovation has occurred. On the one hand, patents restrict access to innovative knowledge (Φ). But on the other hand, they encourage disclosure, leading to subsequent innovation that could not have happened under secrecy ($-sV$). Granting a patent actually increases ex-post welfare ($D < 0$) if $s > \Phi/V$.

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 $b(V - \Phi) + (1 - b)V = V - b\Phi$. Taking the difference between the two, we obtain

$$L = b\Phi. \tag{24}$$

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 (23) and (24), we obtain the following condition for L to be greater than D:

$$s \geq (1 - b)\frac{\Phi}{V}.$$

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 2; we also derive an upper bound on s for secrecy not to be an option ex ante. The profit from research now is $\pi_R(\theta) = \theta - \psi(\theta) + [1 - M(v_{\phi^*})]\phi^*$, while the profit from a bogus patent is $\pi_B(\theta) = b\big[\theta + [1 - M(v_{\phi^*})]\phi^*\big]$. 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 $\psi(0) > (1 - b)[1 - M(v_{\phi^*})]\phi^*$ (replacing property (i) in Assumption B.1).

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 $\theta + [1 - M(v_{\phi^*})]\phi^* - F$ 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. $F = \pi_R(\hat{\theta})$ or $\hat{\theta} + [1 - M(v_{\phi^*})]\phi^* - F = \psi(\hat{\theta})$. 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 B.1, we need $s < 1 - b$. 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.

Appendix C Costly public funds

We now introduce a shadow cost of public funds, $\lambda > 0$. That is, we assume every dollar of funds that needs to be raised through the tax system costs society \$$(1+\lambda)$ due to distortionary taxation. Then, the social-welfare function from (3) becomes

$$\int_{\hat{\theta}}^{\infty} W(\theta)\mathrm{d}G(\theta) - (1-e)L[G(\hat{\theta}) - G(\underline{\theta})] - \gamma(e)[1 - G(\underline{\theta})] - \lambda\big[t - F[1 - G(\underline{\theta})]\big], \tag{25}$$

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 $\alpha = 0$,[31] IR writes

$$t \geq \gamma(e)[1 - G(\underline{\theta})]. \tag{26}$$

Thus, the planner chooses (e, F, t) to maximize (25) subject to (1),(2), (26) 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 C.1. *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

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

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

Since $\partial\ell/\partial F > 0$ by Assumption 2, this simplifies to

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

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.

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.

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.

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.

Dixit, Avinash (2002): Incentives and Organizations in the Public Sector. *Journal of Human Resources*, Vol. 37(4), pp. 696–727.

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.

Prendergast, Canice (2003): The Limits of Bureaucratic Efficiency. *Journal of Political Economy*, Vol. 111(5), pp. 929–958.

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.*