European University Institute

EUI Working Papers

INVENTORS AND IMPOSTORS:
AN ECONOMIC ANALYSIS OF PATENT EXAMINATION

Florian Schuett

EUROPEAN UNIVERSITY INSTITUTE, FLORENCE

MAX WEBER PROGRAMME

Inventors and Impostors:
an Economic Analysis of Patent Examination

FLORIAN SCHUETT

Abstract

The objective of patent examination is to separate the wheat from the chaff. Good applications – those satisfying the patentability criteria, particularly novelty and non-obviousness – should be accepted, while bad applications should be rejected. How should incentives for examiners be designed to further this objective? This paper develops a theoretical model of patent examination to address the question. It argues that examination can be described as a moral-hazard problem followed by an adverse-selection problem: the examiner must be given incentives to exert effort (looking for evidence to reject), but also to truthfully reveal the evidence he finds (or lack thereof). The model can explain the puzzling compensation scheme in use at the U.S. patent office, where examiners are essentially rewarded for granting patents, as well as variation in compensation schemes across patent offices. It also has implications for the retention of examiners and for administrative patent review.

Florian Schuett
Max Weber Fellow 2008-2009, Max Weber Programme, Villa La Fonte, Via delle Fontanelle 10, 50014 San Domenico, Italy
Toulouse School of Economics (GREMAQ).
Email: florian.schuett@eui.eu.

Introduction

Patent examiners at the U.S. Patent and Trademark Office (USPTO) receive a bonus that depends on the number of applications processed. But because a rejection is more time-consuming than a grant, the bonus introduces a bias towards granting patents (Jaffe and Lerner, 2004; Merges, 1999). Such a compensation scheme is puzzling since it does not seem to give examiners good incentives to exert effort. While rejecting an application requires the examiner to come up with evidence that the claimed invention already exists or would have been obvious to someone skilled in the art, granting a patent is easy: the examiner can simply report not having found such evidence. If anything, shouldn't we expect examiners to be rewarded for rejecting applications?

The objective of patent examination is to separate the wheat from the chaff. Good applications – those satisfying the patentability criteria, particularly novelty and non-obviousness – should be accepted, while bad applications should be rejected. How must incentives for examiners be designed to further this objective? In this paper I develop a theoretical model of patent examination to address the question. I argue that examination can be described as a moral-hazard problem followed by an adverse-selection problem: the examiner must be given incentives to exert effort (looking for evidence to reject), but must also be given incentives to truthfully reveal the evidence he finds (or lack thereof). I show that the model can explain the puzzling compensation scheme in use at the USPTO, as well as variation in compensation schemes across patent offices. It also has important policy implications.

In a nutshell, the argument is the following. Suppose the examiner wants to avoid mistakes and believes that a large proportion of applications is bad. If, moreover, he doesn't make much effort in searching for evidence, he will have little confidence that an application is good when the search turns up nothing. Inducing him to truthfully reveal the absence of evidence then requires rewarding him for grants. The possibility of this type of bad equilibrium provides a rationale for the compensation scheme observed at the USPTO.

The argument rests on two premises. First, the signal that an application is bad must be *soft information*, i.e., unverifiable by the principal and third parties. This makes sense because of the technical complexity of patent applications, the vagueness of patentability criteria, and because there is little information on the quality of an examiner's decisions in the short run. While more information becomes available in the long run (e.g., through court decisions on patent validity), this information is difficult to include in a contract. Second, examiners must have a desire to avoid mistakes that is unrelated to short-term monetary compensation. Such a desire might stem from long-term implicit incentives within the organization (promotion etc.), but also from recognition by peers or a concern for social welfare. With a slight abuse of language, I will refer to the desire to avoid mistakes as *intrinsic motivation*.

The argument also raises the question of when a bad equilibrium, characterized by low effort, many bad applications, and – as a result – low-quality patents, will arise. I show that it is precisely when intrinsic motivation is low that such an equilibrium is likely to occur. Low intrinsic motivation leads to lower effort and a larger proportion of bad applications. Under some conditions, ensuring truthfulness then makes it necessary to reward the examiner for grants. As intrinsic motivation increases, however, this may no longer be the case.

I go on to argue that intrinsic motivation is likely to be related to how long the examiner expects to stay at the patent office and to how timely information about the quality of his decisions becomes available. Under this interpretation, the predictions of the model are consistent with casual empirical evidence. A comparison of the USPTO with the European

Patent Office (EPO) shows important differences in examiner turnover, the availability of information on decision quality, short-term compensation, and applicant behavior. The average U.S. examiner stays for only three years while in Europe it is basically a lifetime job (van Pottelsberghe and François, forthcoming). The EPO's opposition system makes information about decision quality available in a more timely manner than court trials, which are the main source of information in the U.S. These facts suggest that intrinsic motivation should be lower in the U.S. than in Europe. Thus, the model would predict that patents issued by the USPTO are of lower quality than EPO patents, and that U.S. examiners are more likely to be rewarded for granting through short-term compensation. At the same time, it makes no prediction on grant rates.

The observation that, unlike their U.S. counterparts, examiners at the EPO receive a fixed wage is in line with these predictions (Friebel *et al.*, 2006). And while patent quality is hard to measure, the perception in the patent community is that the problem is indeed more acute in the U.S.[1] In addition, recent research shows little difference in grant rates between the two offices (Friebel *et al.*, 2006; Lemley and Sampat, 2008).

Why should we care about patent examination? To begin with, patents create (temporary) monopolies. Granting patents for non-inventions causes deadweight loss and litigation without providing any offsetting benefit to society. This would be a minor problem if the courts only enforced good patents. Courts, however, sometimes enforce bad patents, as the near shutdown of BlackBerry in 2006 illustrates.[2] Moreover, many patent disputes never reach the courts. Challenging a bad patent is a public good and may therefore be under-provided (Chiou, 2006; Farrell and Shapiro, 2008). What is particularly troublesome is that, as Chiou (2008) shows, disputes over weak patents are particularly likely to be settled out of court. And when patent disputes do reach the courts, they entail substantial legal costs. Ford *et al.* (2007) estimate the total cost of bad patents to the U.S. economy at an annual $25.5 billion.[3]

In the model presented in section , the government delegates patent examination to an examiner motivated by both extrinsic rewards (i.e., monetary transfers) and intrinsic rewards (defined as a concern about making correct decisions). The examiner must expend effort to obtain a signal about an applicant. If the applicant's claimed invention is not truly new, the examiner can come up with a signal ("prior art") demonstrating the lack of novelty; I assume that the signal is soft information. The examiner takes the proportion of good and bad applications as given. Applicants, however, respond to how rigorous they anticipate examination to be. I assume that the applicants' best-response function is such that the proportion of good applications increases with the expected examination effort.

The government chooses an application fee for firms and an incentive scheme for the examiner. In section , I start by studying the government's choice of incentives, taking the

---

[1] See, e.g., Jaffe and Lerner (2004). The fact that the topic has been much more intensely debated in the U.S. can be seen as a rough indicator that the quality of patents issued by the USPTO is lower.

[2] The maker of BlackBerry mobile devices, Research In Motion (RIM), was sued by patent-holding company NTP, and settled for a reported $612.5 million because the court threatened to issue an injunction unless the parties reached a settlement. The injunction would have shut down BlackBerry. Apparently, the judge was unprepared to wait for the final result of the re-examination of NTP's patents by the USPTO even though the office had indicated that it was likely to revoke 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] Of this sum, they attribute $4.5 billion to litigation costs, while the remainder corresponds to the disincentive to future innovators that patents create. While methodologically controversial, Ford *et al.*'s (2007) calculations indicate that the costs of bad patents are likely to be significant.

application fee as given. Soft information severely limits the use of explicit monetary incentives, so that the examiner's intrinsic motivation becomes the crucial determinant of the equilibrium outcome. I establish two main results. First, both the equilibrium proportion of good applications and the equilibrium effort are increasing in the examiner's intrinsic motivation. Second, for low levels of intrinsic motivation, the optimal incentive scheme rewards the examiner for granting patents. This is true assuming the proportion of bad applications is sufficiently large when applicants expect zero effort. There is a complementarity between intrinsic and extrinsic rewards: the more intrinsically motivated the examiner is, the more effectively can monetary incentives be used. It eventually becomes possible to reward him for rejecting, which feeds back positively into effort provision.

In section , I endogenize the applicants' best-response function. I assume that potential applicants differ in their ability to produce valuable inventions (their creativity) and choose whether to do genuine research or to file applications on existing technologies, hoping to escape detection by the examiner. The profitability of the two activities depends on the examiner's examination effort. More rigorous examination makes it less likely for impostors to obtain patents, and therefore increases the attractiveness of true research. This setup leads to self-selection of applicants. Under a single-crossing condition, high-creativity firms do genuine research, while low-creativity firms submit bad applications or stay idle.

The endogenization of applicant behavior allows me to study the effect of changes in the application fee and to make normative statements about the optimal patent policy. I show that if the government can directly control the level of examination effort, the optimal policy leads to full deterrence of bad applications. Effort is chosen to balance the benefits of research with the costs of patent examination, while the application fee is used to achieve deterrence. When patent examination is delegated to an examiner, however, there is a tradeoff between deterrence and innovation: a lower application fee leads to more bad applications but at the same time induces the examiner to screen more rigorously, which, in turn, leads to more innovation.

In section , I summarize the results of the model and discuss how it relates empirically observed differences between patent offices to compensation schemes and applicant behavior. I also comment briefly on policy implications.

A number of recent papers investigate patent examination. Langinier and Marcoul (2003) and Caillaud and Duchêne (2005) start from the idea that patent examination resembles an inspection game and as such is plagued by commitment problems. 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 who have obtained a positive signal separate from applicants with a negative signal in terms of the amount of prior art they submit. 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. Régibeau and Rockett (2007) examine the optimal duration of patent examination as a function of the importance of an innovation. They find that, controlling for the position in the innovation cycle, more important innovations should be examined faster, a prediction which is born out by evidence from a sample of U.S. patents.

All of these papers consider a benevolent patent office maximizing social welfare. Therefore, they are unable to make predictions about examiner compensation. With the exception

of Caillaud and Duchêne (2005), they also treat the proportion of good and bad applications as exogenous, so they cannot explain differences in applicants' behavior. By contrast, I consider a utility-maximizing examiner (albeit motivated to some extent by a desire to avoid mistakes) and allow the proportion of good applications to depend on the examiner's effort.

The paper is also related to the auditing literature, and particularly Iossa and Legros (2004), who study auditing with soft information. They show that a necessary condition for the auditor to exert any effort is that he be given a stake in the audited project. Similarly, I show that positive effort will only occur if the examiner is intrinsically motivated – that is, if he has a "stake" in the social consequences of his decision.

A simple model of patent examination

Consider the following setup. There are three types of players: a benevolent planner (the government or Congress), a patent examiner, and potential applicants (firms). The planner, whose objective is to maximize social welfare, delegates patent examination to the examiner. Applications filed by firms can be good (G), i.e. true inventions, or bad (B), i.e. non-inventions which already exist or would have been obvious to someone skilled in the art.

Examiner

The examiner does not observe the type of an application but believes that a proportion p is good and a proportion $1 - p$ is bad. He conducts a prior-art search that allows him to receive a signal σ about an application. The distribution of the signal depends on the type of the application and on the examiner's effort, which is unobservable. If the application is good (G), the examiner never obtains any signal ($\sigma = \varnothing$). If the application is bad (B), he obtains a signal $\sigma = B$ with probability e, and no signal with probability $1 - e$, where $e \in [0, 1]$ is the effort that he puts into patent examination.

Assumption 1 (Soft information). *Patent examination produces soft information: the signal $\sigma = B$ is unverifiable by the planner or third parties.*

The examiner has utility
$$U = t + y - \gamma(e),$$

where t is the monetary transfer he receives from the planner, y is an intrinsic reward, and $\gamma(e)$ is the cost of effort (increasing and convex with $\gamma(0) = \gamma'(0) = 0$ and $\gamma'(1) = \infty$). I assume that the examiner is protected by limited liability (i.e., transfers must be non-negative). The intrinsic reward y takes different values depending on the type of application and the approval decision, as indicated in table 1.

| | Application | |
Decision	Good	Bad
Grant	y_G	0
Rejection	0	y_B

Table 1: Intrinsic rewards

Assumption 2 (Intrinsic motivation). *Intrinsic rewards satisfy $y_G \geq 0$ and $y_B \geq 0$.*

According to Assumption 2, the examiner derives an intrinsic reward from accepting good applications and from rejecting bad ones.[4] The *expected* intrinsic reward also depends on the examiner's posterior belief that an application is valid given the result of his prior-art search. This reward structure formalizes the idea that the examiner cares about making the right decision.

Several interpretations are possible. One is that some information about the quality of an examiner's decisions may transpire over time. Although this information cannot be contracted on, it can be used in subjective performance evaluation and thus be brought to bear on promotion and dismissal decisions which are part of the organization's implicit incentives. The information may also be learnt by the examiner's peers, whose esteem he may value. Alternatively, the examiner may have genuine intrinsic motivation, i.e. he may care about the consequences of his decisions on others (in this context, particularly consumers and technology users).[5]

Applicants

Potential applicants' filing strategies depend on how much effort they expect the examiner to provide. For now, I will adopt a reduced-form approach that consists in making assumptions about their best-response function $p(e)$, i.e. the function relating the proportion of good applications to the examiner's effort. In section below, I endogenize applicants' best response by explicitly modeling their filing strategies.

Assumption 3 (Applicants' best response). *Applicants' best-response function $p(e)$ is continuously differentiable and satisfies the following properties: $0 < p(e) \leq 1$ for all e, $p(0) < 1$, and $p' > 0$.*

In words, the proportion of good applicants is always strictly larger than 0 and weakly smaller than 1. When effort is zero, the proportion of bad applications is strictly positive. The proportion of good applications increases with effort.

Timing

The timing of the game is as follows (see figure 1). At the beginning of the game, the planner sets an application fee and chooses an incentive scheme for the examiner.[6] Then, firms file for patents. The examiner 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 examiner cannot commit to a level of examination effort e before firms decide on their filing behavior. This implies that the examiner does not take into account the effect of his effort on the proportion of good and bad applications.

---

[4] The fact that the top-right and lower-left fields are set to zero is a normalization. All that matters for the examiner's decision is the comparison between the intrinsic rewards of granting and rejecting a given type of application.

[5] While the economic literature has only recently begun to acknowledge the importance of intrinsic motivation for understanding bureaucracies (see, e.g., Prendergast (2007)), in the public administration literature the concept of "public-service motivation" has a long tradition, and its relevance is empirically established (Perry and Wise, 1990).

[6] This is a restriction on the set of instruments that the planner has at her disposal. In particular, I impose a uniform application fee instead of conditioning fees on the outcome of the examination.

Figure 1: Timing of the examination game

Discussion of assumptions

The setup we have adopted, with the signal being modeled as soft information and the examiner caring about making correct decisions, calls for some justification. Soft information is generally considered a reasonable description of situations involving complex scientific evidence (see, e.g., Shin, 1998). Patent applications are inherently technical and have increased in complexity over time. Moreover, patentability criteria, and the non-obviousness standard in particular, are often vague, somewhat ill-defined concepts. As noted by Jaffe and Lerner (2004, p. 172), "there is an essentially irreducible aspect of judgment in determining if an invention is truly new. After all, even young Albert Einstein faced challenges while assessing applications (...) in the Swiss Patent Office." In an experiment carried out by the UK Patent Office in 2005, workshop participants were asked to evaluate whether a number of fictitious inventions satisfied different definitions of a "technical contribution" (Friebel *et al.*, 2006).[7] 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.

Moreover, little information about the quality of their decisions is available in the short run. While judicial and administrative review of patent validity, such as court hearings, re-examination (in the U.S.) or opposition (in Europe), provides such information, it occurs with a significant time lag. Another problem is that courts may differ in their "patent friendliness" across time and space.[8] These considerations make it impractical to include information on decision quality in a contract. It seems more appropriate to model it as being part of the implicit incentives within the patent office.[9]

Designing incentives for the examiner

In this section, I look at the design of incentives for the examiner taking the application fee as given. I defer the specification of the social value of examination to section ; for now I assume simply that the examiner's intrinsic motivation alone leads to insufficient effort provision from the planner's point of view. The planner would thus like to increase effort above the "natural"

---

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

[8] Observers have suggested that this was the case in the United States after the creation of a centralized appeals court for patent disputes, the Court of Appeals for the Federal Circuit (CAFC).

[9] It seems inappropriate to treat this as a standard career-concerns setup. The main outside opportunity for patent examiners is employment in law firms. But the value of former patent examiners for patent attorneys comes mainly from their inside knowledge of the patent office, rather than from the particular skills they demonstrated during their stay at the office. As a matter of fact, examiners often leave before any information about the quality of their decisions becomes available to the public. The signaling motive emphasized by career-concerns models seems to be largely irrelevant.

level. I also assume that there is no reason other than (possibly) incentives not to grant a patent when no signal is found.[10]

The problem the planner faces is one of moral hazard followed by adverse selection: the examiner's 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. The planner offers a menu of contracts $(t_{\tilde{\sigma}}, x_{\tilde{\sigma}})$ where $\tilde{\sigma} \in \{B, \varnothing\}$ is the signal reported by the examiner, t is the transfer he receives and x the probability that the patent is granted. That is, the planner asks the examiner to report his signal σ. If he reports B, the planner pays t_B and grants a patent with probability x_B. If he reports $\varnothing$, the planner pays $t_\varnothing$ and grants with probability $x_\varnothing$.

Consider the case where the examiner has exerted equilibrium effort $e^* > 0$ and come up with signal $\sigma = B$. For him to prefer to report B, it must be the case that

$$t_B + (1 - x_B)y_B \geq t_\varnothing + (1 - x_\varnothing)y_B. \tag{1}$$

Given signal B, he knows with certainty that the application is bad, but he only enjoys the intrinsic reward from rejection with probability $(1 - x_{\tilde{\sigma}})$. If, on the other hand, the examiner obtains no signal ($\sigma = \varnothing$), he will prefer to report $\varnothing$ provided

$$t_B + \hat{p}x_B y_G + (1 - \hat{p})(1 - x_B)y_B \leq t_\varnothing + \hat{p}x_\varnothing y_G + (1 - \hat{p})(1 - x_\varnothing)y_B, \tag{2}$$

where $\hat{p} \equiv \Pr[G|\varnothing]$ is the examiner's posterior belief that the application is valid given that he has found no evidence to the contrary. His expected intrinsic reward from reporting B is $\hat{p}x_B y_G + (1 - \hat{p})(1 - x_B)y_B$, while that from reporting $\varnothing$ is $\hat{p}x_\varnothing y_G + (1 - \hat{p})(1 - x_\varnothing)y_B$.

Turning to the moral-hazard stage, suppose the examiner anticipates truthfully revealing the signal he finds. He then chooses e to maximize

$$p[t_\varnothing + x_\varnothing y_G] + (1 - p)\big[e[t_B + (1 - x_B)y_B] + (1 - e)[t_\varnothing + (1 - x_\varnothing)y_B]\big] - \gamma(e).$$

With probability p, the application is good, so that he cannot find any grounds for rejection. The transfer he receives is $t_\varnothing$, and the expected intrinsic reward is $x_\varnothing y_G$. With probability $1 - p$, the application is bad, for which he finds evidence with probability e. He is paid t_B and enjoys an expected intrinsic reward of $(1 - x_B)y_B$. With probability $1 - e$, the examiner finds no evidence. He receives a transfer of $t_\varnothing$ and an expected intrinsic reward of $(1 - x_\varnothing)y_B$. Differentiating with respect to e leads to the first-order condition

$$(1 - p)[t_B - t_\varnothing - (x_B - x_\varnothing)y_B] = \gamma'(e). \tag{3}$$

It follows from (3) that effort is increasing in $t_B - t_\varnothing$ and decreasing in $x_B - x_\varnothing$. Moreover, a strictly positive level of examination effort is only sustainable if the examiner expects there to be some bad applications ($p < 1$).

A final set of constraints comes from the possibility of double deviation: the examiner may deviate from both the equilibrium effort and truthful reporting. Two cases are relevant: always reporting B, and always reporting $\varnothing$.[11] In both cases, choosing $e = 0$ is optimal (if

---

[10] I discuss this assumption in footnote 15 below.

[11] A third strategy, which would consist in always reporting the opposite of the signal found, leads to an optimal effort of zero and therefore reduces to the strategy of always reporting B.

the examiner anticipates that his report will not depend on his signal, there is no point in exerting effort). To rule out double deviation, the equilibrium utility with truthful reporting must be larger than the utility with zero effort and either report (B or $\varnothing$). Letting U^* denote the examiner's equilibrium utility, we must have

$$t_B + p x_B y_G + (1-p)(1-x_B)y_B \leq U^* \tag{4}$$

and

$$t_\varnothing + p x_\varnothing y_G + (1-p)(1-x_\varnothing)y_B \leq U^*, \tag{5}$$

with

$$U^* = p[t_\varnothing + x_\varnothing y_G] + (1-p)\big[e^*[t_B + (1-x_B)y_B] + (1-e^*)[t_\varnothing + (1-x_\varnothing)y_B]\big] - \gamma(e^*). \tag{6}$$

Given the absence of a shadow cost of public funds, transfers are not costly to the planner (they are pure redistribution). Moreover, as I show in the proof of Lemma 1 below, there is no conflict between welfare and effort maximization in the choice of grant probabilities. Therefore, the planner's objective is simply to maximize the examiner's effort. Since incentives for effort provision are increasing in the left-hand side of (3), the planner's problem is

$$\max_{(t_B, x_B), (t_\varnothing, x_\varnothing)} t_B - t_\varnothing - (x_B - x_\varnothing)y_B,$$

subject to (1), (2), (4), (5), $t_\sigma \geq 0$, $0 \leq x_\sigma \leq 1$, and $e^* \leq e^o$, where e^o denotes the level of effort the planner would choose if he could control it directly. Ignoring the last constraint, we have:

Lemma 1 (Incentive design). *In designing the examiner's incentives, the planner optimally chooses deterministic grant probabilities: $x_\varnothing = 1$ and $x_B = 0$. The optimal transfers satisfy*

$$t_B - t_\varnothing = \hat{p}y_G - (1-\hat{p})y_B - \frac{\gamma(e^*)}{p + (1-p)(1-e^*)}. \tag{7}$$

Proof: Since incentives for effort provision increase with $t_B - t_\varnothing$, it is the upward constraints, (2) and (4), that are relevant. Rewriting them respectively as

$$t_B - t_\varnothing \;\leq\; (x_\varnothing - x_B)[\hat{p}y_G - (1-\hat{p})y_B] \tag{8}$$

$$(t_B - t_\varnothing)[p + (1-p)(1-e^*)] \;\leq\; (x_\varnothing - x_B)[p y_G - (1-p)(1-e^*)y_B] - \gamma(e^*) \tag{9}$$

and using $\hat{p} = p/[p + (1-p)(1-e)]$ so that (9) becomes

$$t_B - t_\varnothing \leq (x_\varnothing - x_B)[\hat{p}y_G - (1-\hat{p})y_B] - \frac{\gamma(e^*)}{p + (1-p)(1-e^*)}, \tag{10}$$

we see that (10) implies (8). Thus, (10) is the binding constraint, which we can use to replace $t_B - t_\varnothing$ in the objective. We obtain

$$(x_\varnothing - x_B)\hat{p}[y_G + y_B - \gamma(e^*)/p]. \tag{11}$$

Notice that any incentive-compatible contract will feature $x_\varnothing \geq x_B$; otherwise (1) and (2) cannot be simultaneously satisfied. Thus, if there is to be a positive level of effort, the

Figure 2: Equilibrium of the examination game

expression in parentheses must be positive. It follows that (11) is increasing in $x_\varnothing - x_B$, so incentives are maximized for $x_\varnothing = 1$ and $x_B = 0$. Substituting these values in (10) yields the claimed result. ∎

Lemma 1 shows that, in terms of incentives, applications should always be rejected when defeating prior art is found, and granted when none is found. The intuition is that, even though reducing $x_\varnothing$ or increasing x_B can relax the incentive-compatibility constraints (by making lying less tempting), it also weakens the incentive to provide effort. The second effect dominates, making it optimal to use deterministic grant probabilities. The difference in transfers, $t_B - t_\varnothing$, is chosen at the highest level compatible with the double-deviation constraint (4).[12]

Incentive compatibility severely limits the use of monetary transfers by imposing an upper bound on the power of incentives, as equation (7) shows. The intuition is that soft information gives the examiner discretion over the grant decision. If we pay him a lot for rejecting, he will reject too many applications (in this simple model, all of them in fact). If we pay him a lot for accepting, he will accept too many of them. If he is to exert any effort, he must anticipate truthfully revealing the signal he finds. Monetary incentives can only induce additional effort to the extent that they do not jeopardize truthful revelation.

Having derived the optimal incentive scheme, we can compute the examiner's best-response function $e(p)$, i.e., the function that relates his effort to the proportion of good applications. Plugging the values from Lemma 1 into (3), $e(p)$ is obtained as the solution to

$$(1-p)\frac{p[y_G + y_B] - \gamma(e)}{p + (1-p)(1-e)} = \gamma'(e). \tag{12}$$

Combining applicants' and the examiner's best responses yields the equilibrium of the examination game.

---

[12] Note that the lemma does not specify the level of transfers, but only the difference. The level will be chosen so as to satisfy the examiner's participation constraint, which I have not made explicit because public funds are assumed to be costless.

Lemma 2 (Existence of equilibrium)**.** *There exists an equilibrium (p^*, e^*) of the examination game characterized by*

$$p^* = p(e^*)$$
$$\gamma'(e^*) = (1 - p^*)\frac{p^*[y_G + y_B] - \gamma(e^*)}{p^* + (1 - p^*)(1 - e^*)}.$$

Proof: By Assumption 3 and equation (12), both $p(e)$ and $e(p)$ are continuous. Also by Assumption 3, $p(e)$ is monotone increasing and defined on $[0, 1] \rightarrow [p(0), p(1)]$, where $0 < p(0) < p(1) \leq 1$. Furthermore, since $p = 0$ and $p = 1$ are solutions to (12) when e is set to zero, it follows that $e = 0$ is part of the best-response correspondence $e(p)$ for $p = 0$ and $p = 1$. By continuity, $e(p)$ is defined on $[0, 1] \rightarrow [0, e_{\max}]$, where $e_{\max} \equiv \max_p e(p)$. Hence, equilibrium exists. ∎

The equilibrium is depicted in figure 2. The inverted-U shape of the $e(p)$ function has an intuitive explanation. If $p = 0$ or $p = 1$, the examiner knows in advance whether he is facing a good or bad application. There is no point in exerting effort to acquire information that is redundant. Lemma 2 shows existence of equilibrium. The assumptions made do not guarantee uniqueness, however. What is important for the remainder of the analysis is that in the case of multiple equilibria the equilibrium is picked according to a deterministic rule rather than randomly.

The restrictions on transfers caused by soft information (see Lemma 1) mean that extrinsic rewards can only play a limited role in providing incentives. This gives a crucial role to intrinsic motivation. In the following proposition, I introduce a constant α, by which I multiply both types of intrinsic reward, y_G and y_B; α can be interpreted as a measure of the overall strength of intrinsic motivation, keeping the ratio between y_G and y_B fixed. It allows us to analyze how intrinsic motivation affects the equilibrium outcome, all other things being equal. The idea is that the relative strength of y_G and y_B is largely determined exogenously, for example by applicants' and challengers' propensities to appeal the examiner's decisions. The absolute strength of intrinsic motivation is likely to be more malleable to policy intervention.

Proposition 1 (Importance of intrinsic motivation)**.** *Let $\alpha \geq 0$ be a constant multiplying y_G and y_B. If $\alpha = 0$, no effort can be sustained in equilibrium. An increase in α leads to an equilibrium with greater effort and a larger proportion of good applicants.*

Proof: The first part of the proposition is immediate from looking at equation (12), determining $e(p)$; if $\alpha = 0$, then $e(p) = 0$ for all p. For the second part, it suffices to show that $\mathrm{d}e(p)/\mathrm{d}\alpha \geq 0$ because, by Assumption 3, $p' > 0$. Rewrite (12) as

$$(1 - p)p\alpha(y_G + y_B) = \gamma'(e)[p + (1 - p)(1 - e)] + (1 - p)\gamma(e).$$

By the implicit function theorem, $\mathrm{d}e(p)/\mathrm{d}\alpha$ has the sign of the derivative of the right-hand side with respect to e. Computation yields

$$\frac{\partial}{\partial e}[\gamma'(e)[p + (1 - p)(1 - e)] + (1 - p)\gamma(e)] = \gamma''(e)[p + (1 - p)(1 - e)] \geq 0,$$

where the inequality follows from the convexity of γ. ∎

Figure 3: Effect of an exogenous increase in intrinsic motivation

Some amount of intrinsic motivation is essential for effort provision. If $\alpha = 0$, the examiner responds to any p with zero effort. Proposition 1 is illustrated in figure 3, which depicts the effect of an exogenous increase in intrinsic motivation from $\alpha > 0$ to $\alpha' > \alpha$. An examiner who cares more about making the right decision exerts more effort, whatever the proportion of good and bad applications. The $e(p)$ function shifts out, and the equilibrium moves to the north-east, along the $p(e)$ curve.

The second main result of the equilibrium analysis concerns the compensation scheme, and is the subject of Proposition 2. The constant α again measures the strength of intrinsic motivation.

Proposition 2 (Examiner compensation)**.** *Suppose $p(0) < y_B/(y_B + y_G)$ and suppose that there is $\tilde{e} < 1$ such that $p(\tilde{e}) = 1$ and $\gamma(\tilde{e}) < y_G$. Then, there exists a threshold $\hat{\alpha} \in (0, \infty)$, such that $t_B < t_\varnothing$ for $0 < \alpha < \hat{\alpha}$ and $t_B > t_\varnothing$ for $\alpha > \hat{\alpha}$.*

Proof: Suppose $\alpha = 0$. By Proposition 1, we then have $e^* = 0$ and $p^* = p(0)$. From Lemma 1, it follows that $t_B - t_\varnothing = 0$. Compute

$$\frac{\mathrm{d}(t_B - t_\varnothing)}{\mathrm{d}\alpha} = \frac{\mathrm{d}\hat{p}}{\mathrm{d}\alpha}\alpha[y_G + y_B] + \hat{p}y_G - (1 - \hat{p})y_B - \frac{\hat{p}}{p}\frac{\mathrm{d}e}{\mathrm{d}\alpha}\gamma'(e) - \frac{\gamma(e)}{p^2}\left[\frac{\mathrm{d}\hat{p}}{\mathrm{d}\alpha}p - \frac{\mathrm{d}p}{\mathrm{d}\alpha}\hat{p}\right].$$

Evaluating this expression at $\alpha = 0$, noting that $\hat{p} = p$ for $e = 0$, we obtain

$$\left.\frac{\mathrm{d}(t_B - t_\varnothing)}{\mathrm{d}\alpha}\right|_{\alpha=0} = p(0)y_G - (1 - p(0))y_B < 0.$$

It follows that for small values of α, $t_B - t_\varnothing$ is negative. As α increases, so do e^* and p^*, by Proposition 1. Eventually, $e^* \to \tilde{e}$ and, by the definition of $\tilde{e}$, $p^* \to 1$, also implying $\hat{p} = 1$. Thus, $t_B - t_\varnothing \to y_G - \gamma(\tilde{e})$, which is positive by assumption. It follows that there must exist a threshold $\hat{\alpha}$ as defined in the proposition. ∎

Proposition 2 says that if intrinsic motivation is low (and under some conditions on $p(e)$), the compensation scheme rewards the examiner for granting. Such a reward is needed to

11

ensure truthful revelation: were he not compensated for granting by means of a monetary transfer, the examiner would reject all applications. The intuition is that in an equilibrium where the proportion of bad applications is large and effort low, the best the examiner can do to avoid mistakes is reject everything. Of course, anticipating this, he will not exert any effort either. Thus, the reward for granting actually induces positive, albeit low, effort.

As intrinsic motivation increases, it eventually becomes possible to reward the examiner for rejecting applications without impeding truthful revelation. Rewarding rejection has a positive feedback effect on effort. The model thus yields a complementarity between intrinsic and extrinsic rewards: higher intrinsic motivation allows the planner to use monetary incentives more effectively. Intuitively, as the equilibrium values of p and e increase, the conflict between truthful revelation and effort provision is attenuated.

Endogenizing applicant behavior

Modeling firms' choice of activity

In this section, I endogenize the applicants' best-response function $p(e)$. This allows me to derive some further results relating to the planner's choice of the application fee ϕ. Suppose there is a continuum (with mass 1) of potential applicant firms. Firms are characterized by a creativity parameter θ, which is their private knowledge and distributed according to cdf $F(\cdot)$ on $[0, \infty)$.

Assumption 4 (MHRP). *The distribution of θ satisfies the monotone hazard rate property:*

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

Firms are endowed with one indivisible unit of time which they can devote either to R&D or to filing a bogus patent application claiming something that is either obvious or not novel. Alternatively, firms can stay idle. The idea is that there are existing technologies or obvious combinations of existing technologies that (a) firms can claim to have invented and which are not easily distinguishable from true inventions, and that (b), if awarded a patent, allow the patent holder to extract rents from users; a necessary condition is that such bad patents are enforced by the courts with positive probability. Denote a firm's decision by $d(\theta) \in \{R, B, I\}$. If it does R&D ($d(\theta) = R$), its payoff when awarded a patent is $\pi_R(\theta)$.[13] 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 the possibility that the patent may be invalidated by the courts later on). I assume that firms' profit is zero or negative if they fail to obtain a patent. Their payoff when staying idle ($d(\theta) = I$) is zero.

Given an application fee ϕ and an anticipated examination effort e, each type of firm chooses $d(\theta)$ to maximize its expected payoff. Suppose for simplicity that research always leads to patentable inventions. Research then yields a net profit of $\pi_R(\theta) - \phi$, while a bogus applicant can expect net profit $(1-e)\pi_B(\theta) - \phi$.[14] Thus, a firm prefers R&D to imposture if

---

[13] This can be seen as a reduced-form profit function resulting from a firm's investment choice; see footnote 14 below.

[14] The assumption that genuine research always results in patentable inventions is not crucial. If instead genuine inventors sometimes inadvertently re-invent old products or processes, their expected profit decreases with e. But what matters for the decision between research and imposture is the relative attractiveness of

and only if

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

Assumption 5 (Single crossing). *Profit functions satisfy*

(i) $\pi_R' > \pi_B' > 0$,

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

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

(iv) $\pi_R'' \leq 0$ *and* $\pi_B'' \geq \pi_R''$.

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

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

$$\pi_R(\hat{\theta}) = (1 - e)\pi_B(\hat{\theta}). \tag{13}$$

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

$$(1 - e)\pi_B(\underline{\theta}) = \phi, \tag{14}$$

such that firms with creativity higher than $\hat{\theta}$ do research, firms with creativity between $\hat{\theta}$ and $\underline{\theta}$ submit bogus applications, and firms with creativity lower than $\underline{\theta}$ remain idle. Thus, a patent policy (ϕ, e) leads to self-selection of firms between genuine R&D, imposture, and inactivity, as illustrated in figure 4.

To close the model, I specify the effects of innovations and bad patents on social welfare. Innovations generate social welfare (profits plus consumer surplus) $W(\theta) \geq \pi(\theta)$ with $W' \geq 0$. That is, the social value of innovation (weakly) exceeds the private value, and more creative inventors produce more valuable innovations, both from a private and a social point of view. Bad patents cause a social loss of $L(\theta) > 0$. I make no assumption on how this loss is related to creativity.

---

each of these activities. Increasing e still makes research relatively more attractive than imposture.

If $\pi_R(\theta)$ is interpreted as a reduced-form profit function resulting from the firm's investment choice, another question is whether examination effort and the 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 ϕ. To see this, assume (following Cornelli and Schankerman (1999)) that the firm's profit (gross of application fees) is given by $\rho(z, \theta) - \psi(z)$, where z is its R&D investment and $\psi(z)$ the associated cost. Assuming $\rho_z > 0 \geq \rho_{zz}$ (subscripts denote partial derivatives), and $\psi' > 0$, $\psi'' > 0$, 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 ϕ, and $\pi_R(\theta) = \rho(z^*(\theta), \theta) - \psi(z^*(\theta))$.

Figure 4: Self-selection of firms according to creativity θ

Optimal policy when the planner directly controls examination

Let us derive the patent policy that the planner would choose ex ante if she could directly control both application fee and effort.[15] The optimal combination of ϕ and e maximizes

$$\int_{\hat{\theta}}^{\infty} W(\theta) \mathrm{d}F(\theta) - (1-e) \int_{\underline{\theta}}^{\hat{\theta}} L(\theta) \mathrm{d}F(\theta) - \gamma(e)[1 - F(\underline{\theta})] \tag{15}$$

subject to (13), (14) and $\underline{\theta} \leq \hat{\theta}$. The first term corresponds to the social value created by research (undertaken by firms whose creativity exceeds $\hat{\theta}$), the second term captures the expected social losses from bad patents, and the third term represents the cost of examination. The constraint $\underline{\theta} \leq \hat{\theta}$ reflects the fact that setting e and ϕ such that $\hat{\theta}$ is strictly below $\underline{\theta}$ can never be optimal. Holding ϕ 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 3 (Optimal policy). *Suppose Assumptions 4 and 5 hold. The optimal policy (e^o, ϕ^o) involves full deterrence of bad applications: $\underline{\theta} = \hat{\theta}$. Examination effort e^o satisfies the following equation:*

$$-h'(e^o)W(\hat{\theta})f(\hat{\theta}) = \gamma'(e^o)[1 - F(\hat{\theta})] - h'(e^o)\gamma(e^o)f(\hat{\theta}). \tag{16}$$

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

Proof: Let us first show that the constraint $\underline{\theta} \leq \hat{\theta}$ must be binding. Let μ be the multiplier associated with the constraint. Differentiating (15) with respect to ϕ, we have

$$\frac{\partial \ell}{\partial \phi} \left[f(\underline{\theta})[(1-e)L(\underline{\theta}) + \gamma(e)] - \mu \right] = 0. \tag{17}$$

---

[15] I restrict attention to deterministic grant probabilities, i.e. $x_B = 0$ and $x_\varnothing = 1$. While $x_B = 0$ is clearly optimal, $x_\varnothing = 1$ may not be: by not always issuing a patent when no signal is found, the planner avoids deadweight loss. Because I have not explicitly modeled R&D investment, however, I cannot make a meaningful statement on the optimal $x_\varnothing$ within this model. I therefore assume that $x_\varnothing$ is constrained to be 1 by law.

Since $\partial \ell / \partial \phi > 0$, $\mu > 0$, so indeed $\underline{\theta} = \hat{\theta}$. This implies $\phi = \pi_R(h(e))$. We obtain (16) by differentiating (15) with respect to e, substituting for μ from (17) 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

$$-h''Wf - (h')^2[W'f + Wf'] - \gamma''(1-F) + 2h'\gamma'f + \gamma\left[h''f + h'f'\right] < 0.$$

At e^o, this can be rewritten using the fact that, by (16), $\gamma = \gamma'(1-F)/(h'f) + W$:

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

The fraction is positive thanks to Assumption 4. 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 5. ∎

    Greater examination effort increases the attractiveness of genuine research relative to imposture. That is, e determines the incentives to do R&D. The planner chooses e^o to equalize the marginal social gains from more innovation (the left-hand side of (16)) with the marginal cost of examination (the right-hand side of (16)). Meanwhile, ϕ^o is set so as to deter all firms with $\theta < \hat{\theta} = h(e^o)$ from applying. At the optimum, there are no bogus applications, and no bad patent is issued. Intuitively, as long as $\phi < \pi_R(\hat{\theta})$, raising the application fee does not represent a disincentive to innovation in this model: only those types of firm who would anyway find it optimal to submit bogus applications are discouraged from applying for patents. Thus, there is no loss in raising the fee up to the level where imposture is completely deterred.

Choice of application fee with delegated examination

The previous section analyzed the benchmark case where the planner directly controls e. I now return to the case where examination is delegated to an examiner and investigate the planner's choice of application fee when she cannot control e directly but only indirectly through the examiner's incentive scheme. I start by showing that the best-response function generated by the model of applicant behavior in section satisfies Assumption 3, so the results from section continue to apply. I then investigate the effect of the application fee ϕ on the applicants' best response and draw some conclusions for the planner's choice of ϕ.

    Firms' best response to examination effort e is

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

Figure 5: Effect of a change in the application fee

Since activity R always results in patentable inventions, and thus good applications, while activity B always results in bad ones, the thresholds $\hat{\theta} = h(e)$ and $\underline{\theta} = \ell(e, \phi)$ determine the proportion of good applications:

$$p(e) = \frac{1 - F(h(e))}{1 - F(\ell(e, \phi))}.$$

Since F, h and ℓ are continuously differentiable functions, so is p. It is bounded below by $p(0) = \frac{1 - G(h(0))}{1 - G(\ell(0, \phi))} > 0$ (by property (iii) of Assumption 5) and bounded above by 1. For a given ϕ, the upper bound is reached at $\tilde{e}$, defined by $h(\tilde{e}) = \ell(\tilde{e}, \phi)$. Since $\mathrm{d}h/\mathrm{d}e < 0$ and $\partial \ell / \partial e > 0$, we have $p'(e) > 0$. Thus, the model of endogenous applicant behavior satisfies Assumption 3.

How does the application fee affect the applicants' best response? Since $\partial \ell / \partial \phi > 0$, we have $\partial p(e) / \partial \phi \geq 0$ for all e. Thus, the proportion of good applications is increasing in the application fee, as depicted in figure 5.

Let us consider the comparative statics of a change in the application fee. As ϕ increases, the $p(e)$ curve shifts upwards. 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 small), raising ϕ leads to increases in both p^* and e^*. In the downward-sloping part (p large), raising ϕ leads to an increase in p^* and a decrease in e^*.

It follows that it can never be optimal for the planner to choose ϕ such that the equilibrium is in the upward-sloping part of the $e(p)$ curve; increasing ϕ 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 bad applications. On the other hand, the resulting decrease in equilibrium effort reduces the level of innovation. 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.

Conclusion

I have presented a three-tier hierarchy model of patent examination. A benevolent planner (the principal) delegates patent examination to an examiner (the supervisor) who receives applications filed by firms (the agents). The planner chooses an application fee for firms and an incentive scheme for the examiner. An application can be good or bad, and the examiner needs to exert effort to obtain a signal about it. I model examination as a moral-hazard problem followed by an adverse-selection problem: the examiner must be induced to provide effort but also to reveal the signal he finds, the assumption being that the signal is *soft information* (unverifiable by third parties, including the planner). I have also assumed that the examiner has a desire to make the right decisions, which I have termed *intrinsic motivation*. Finally, I have modeled the proportion of good applications as endogenous, depending on the effort that firms expect the examiner to provide.

I have shown that soft information severely constrains the design of incentives, so that intrinsic motivation becomes a crucial determinant of the equilibrium outcome. When intrinsic motivation is low, the equilibrium features low effort and a large proportion of bad applications. In such an equilibrium, monetary incentives may be reduced to the role of ensuring truthful revelation, leading to a seemingly paradoxical compensation scheme that rewards examiners for granting. Yet this scheme succeeds in inducing the examiner to provide effort: if the examiner anticipated not being truthful, he would optimally choose zero effort. The model also generates a complementarity between intrinsic and extrinsic rewards. As intrinsic motivation increases, extrinsic (monetary) incentives can be used more effectively.

I have argued that the modeling assumptions I use (most notably soft information and intrinsic motivation) provide a reasonable description of how patent examination works in practice. Examining patents requires assessing complex scientific evidence. Moreover, there is little short-term information about the quality of the examiner's decisions; such information only becomes available after a delay and is difficult to contract on. It may, however, be used in the organization's promotion and dismissal decisions, which provide long-term implicit incentives. These implicit incentives tend to create a desire to make correct decisions on the examiner's part, consistent with how I have defined intrinsic motivation.

If the examiner cares about correct decisions because they affect his future with the patent office, a case can be made that how much he cares depends on how long he expects to stay at the patent office. He is likely to care more if he expects to stay long-term because, in the long run, more information about the quality of his decision-making becomes available. He can be rewarded for good decisions through promotion and punished for poor decisions through dismissal. For the same reason, intrinsic motivation is also likely to depend on the precise meaning of "long run." That is, how timely does information about the examiner's decisions become available?

On both of those dimensions, the U.S. and European patent offices differ considerably. At the EPO, examiners usually stay for their entire career. At the USPTO, the average examiner stays for only three years, making long-term incentives largely irrelevant. The EPO also has the edge in terms of timely information about decision quality, thanks to its widely-used opposition system. Opposition allows private parties to mount a challenge against questionable patents through the patent office itself. The opposition procedure produces much faster results than judicial review through the court system. Although the USPTO has a similar procedure called re-examination, it is rarely used (Graham *et al.*, 2002).

In the light of these considerations, which suggest that intrinsic motivation, as defined in

this paper, is higher at the EPO than at the USPTO, the model can explain why U.S. examiners are essentially rewarded for granting patents, but also why European examiners do not face a similar compensation scheme and instead receive a fixed wage. In addition, its predictions are consistent with the fact that the quality of patents issued is generally perceived to be lower in the U.S. than in Europe.

The main policy implications concern examiner retention and administrative patent review. A functioning system of administrative review makes information on the examiners' decision quality available in a more timely manner. Examiners should be retained long enough for long-term incentives to be effective. While this probably requires increasing their salary to match their outside opportunity, the resulting improvement in the quality of examination should reduce the number of bad applications filed. This will partially offset the effect of increasing salaries on costs.

The analysis suggests that retaining examiners and creating administrative review are important for reasons beyond those typically mentioned in the patent-reform debate, which has focused on the fact that more experienced examiners perform better work and that private parties may be better informed about prior art than examiners. Rather, the argument here is that both measures improve examiners' incentives to make correct decisions and allow for more effective reinforcement of effort provision through short-term compensation.

References

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

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

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

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

Farrell, Joseph and Shapiro, Carl (2008): How Strong Are Weak Patents? *American Economic Review*, Vol. 98(4), pp. 1347–1369.

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.

Graham, Stuart J., Hall, Bronwyn H., Harhoff, Dietmar, and Mowery, David C. (2002): Post-Issue Patent 'Quality Control': A Comparative Study of US Patent Re-examinations and European Patent Oppositions. NBER Working Paper no. 8807.

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.

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

Lemley, Mark A. and Sampat, Bhaven N. (2008): Is The Patent Office A Rubbers Stamp? *Emory Law Journal*, Vol. 58, pp. 101–128.

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

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

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

Régibeau, Pierre and Rockett, Katharine (2007): Are More Important Patents Approved More Slowly and Should They Be? CEPR Discussion Paper no. 6178.

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

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