The Benefit of the Doubt: Patent Examination under Strategic Obfuscation*

Nicolás Figueroa† and Jorge Lemus‡

January 5, 2023

Abstract

An applicant privately observes the value of an idea and files a patent application. After the examination, the patent examiner may not find conclusive evidence of the idea's value. In this case, granting a patent can be the right decision or a mistake. Moreover, it can prompt *obfuscation*—effort to disrupt communication—by the applicant, reducing allocative efficiency. We show there is more obfuscation when examiners are more lenient, and provide an empirical exercise supporting this conclusion. We also show that more examination resources do not prevent and may lead to more obfuscation. This result continues to hold when applicants can invest in knowledge to increase their chances of having good ideas. We also show that there is less obfuscation when good ideas are harder to find.

Keywords: Obfuscation, strategic communication, patents, patent examination
JEL classifications: O3, D8, H8

*For comments and suggestions, we thank seminar participants at the University of Connecticut, Universidad Católica, Michigan State University, Hong-Kong University of Science and Technology, and Cornell (Johnson), and LACEA (2022).

†Pontificia Universidad Católica, Chile. Email: nicolasf@uc.cl

‡University of Illinois at Urbana-Champaign. Email: jalemus@illinois.edu

1 Introduction

Understanding the value of an idea is crucial to allocating resources efficiently. Communication, however, is imperfect, so a receiver may not understand the idea communicated by a sender. When the receivers' goal is to approve good ideas only, what should they do with ideas they do not understand? By rejecting all of them, receivers err on the side of caution, rejecting some good ideas; By approving all of them, receivers approve some bad ideas, contravening their goal. Moreover, by giving the benefit of the doubt, a lenient receiver can motivate senders with bad ideas to strategically "obfuscate," i.e., purposely exerting effort to disrupt communication.

Our paper centers on strategic obfuscation by patent applicants, although obfuscation arises in any setting where conclusive evidence lowers the sender's payoff. In patent examination, there is conflict of interest between applicants and examiners: applicants would like all their inventions patented, whereas the examiners' goal is to grant patents only to patentable ideas. Examiners reject an application when they find evidence indicating it is not patentable, i.e., it is not novel, not useful, or obvious (Atal and Bar, 2010). An examiner who remains *uncertain* about the idea's value after the examination of an application may give the benefit of the doubt and grant a patent. If that is the case, an applicant aware of some piece of evidence that would trigger rejection has incentives to hinder the examiner's ability to find it, i.e., the applicant *obfuscates* the application.

Obfuscation has been empirically documented in patent examination (see, e.g., Freilich, 2017; Ashtor, 2022), and it can take many forms: adding unnecessary information, using strategic citations, or writing in complicated language. A lab experiment by Jin et al. (2021) finds that senders strategically obfuscate and receivers make systematic mistakes. In practice, examiners are often unable to identify obfuscation strategies, at least in the short run. Further, if one obfuscation strategy (e.g., lengthy application) becomes pervasive, examiners would adjust their examination strategy, rendering that obfuscation strategy ineffective. Strategic applicants would find another, perhaps subtler one (e.g., a shorter application with more complicated words), such that the examiner does not recognize this new strategy as obfuscation.[1]

In our setting, there are two players: an applicant and an examiner. The applicant privately learns the value of her idea (good or bad), chooses an obfuscation effort, and files a patent

[1] For instance, following the development of textual-analysis algorithms, firms prepare their financial reports to avoid cues perceived as negative by algorithms (Cao et al., 2020).

application. She gets a positive payoff if the examiner grants a patent and zero otherwise. The examiner reviews the application, trying to learn the idea's value. The examiner does not observe the applicant's obfuscation effort, and there are three outcomes of the examination process: (1) hard evidence that the idea is good; (2) hard evidence that the idea is bad; (3) the examiner remains *uncertain* about the idea's value after the examination. Finally, after observing the outcome of the examination process, the examiner chooses to grant a patent or reject the application. The examiner receives a positive payoff from patenting a good idea and incurs a cost from patenting a bad one.

The examiner's decision is trivial whenever the examination reveals the idea's value: she patents good ideas and rejects bad ones. If the examination does not reveal the idea's value, we say that the examiner is *uncertain*. Applicants with good ideas want the examiner to learn the value of their ideas, so they never obfuscate. In contrast, applicants with bad ideas do not want the examiner to learn that their idea is bad because they may get a patent when the examiner is uncertain. As such, applicants with bad ideas might obfuscate depending on how much they value receiving a patent, their obfuscation technology, the examiner's resources, and, crucially, whether uncertain examiners grant patents, which is an equilibrium decision.

An uncertain examiner never grants a patent when a priori the expected benefit of patenting is larger than the expected cost (e.g., when good ideas are scarce). An examiner is more pessimistic (relative to the prior) about an idea's value when they are uncertain since only applicants with bad ideas obfuscate. If uncertain examiners never grant patents, obfuscation does not create any benefit, and it is costly, hence, no applicant obfuscates.

By contrast, when a priori the expected benefit of patenting is *much* larger than the expected cost (e.g., when good ideas are very abundant, very valuable, or the cost of patenting a bad idea is low) uncertain examiners are *lenient* and always grants a patent because patenting bad ideas is a risk worth taking for them. This motivates applicants with bad ideas to obfuscate in an attempt to make the examiner uncertain.

When the expected benefit of patenting is slightly larger than the expected cost (e.g., when good ideas are not too scarce nor too ubiquitous), an uncertain examiner cannot grant a patent, or reject an application, with certainty in equilibrium. If uncertain examiners granted patents with certainty, applicants with bad ideas would obfuscate too much, so the uncertain examiners would prefer to reject the application. And if they rejected applications with certainty, applicants with bad ideas would not obfuscate, so uncertain examiners would prefer to grant a patent for sure since ideas are a priori worth funding. Thus, in equilibrium,

applicants must obfuscate up to the point of leaving an uncertain examiner indifferent between granting a patent and rejecting the application.

One of the predictions of our model is that applicants obfuscate more in environments where a priori examination is more lenient, i.e., when uncertain examiners are more likely to grant patents. We collect data on U.S. patent applications to empirically test this prediction. First, we construct a proxy measure of leniency for each art unit. Some art units are "tougher" and other more "lenient." For instance, in 2015, over 90 percent of applications in "Digital and Optical Communications" received a patent, whereas under 20 percent in "Financial and Business practices." We follow the literature and construct a number of textual measures related to obfuscation in patent applications, including measures of obfuscation in the application claims (indefiniteness, inclusiveness, number of words) and in the description of the invention (concreteness, inclusiveness, number of words). Our results suggest that applicants filing applications in art units with higher patenting rates obfuscate more: claims are more indefinite, inclusive, and have more words; and the description of the invention is less concrete, more inclusive, and have less words. These results together align with our model predictions in that that obfuscation is more common in more lenient art units.

Recent work suggests that patent examiners spend too little time reviewing applications (see, e.g., Frakes and Wasserman, 2017), so they need more resources. We use our results to shed light on policies that increase examination resources, showing that these policies could backfire. The U.S. Patent and Trademark Office recently implemented a policy along these lines.[2] All else equal, examiners with more resources find the value of ideas more frequently, so they reject bad ideas more frequently. Then, strategic applicants with bad ideas obfuscate *more* to make more resourceful examiners uncertain. This strategic response can undo the impact of more resources on improving examination. We show that the examiner's payoff is, in general, non-monotone in the level of resources. More resources benefit the examiner whenever uncertain examiners do not grant patents for sure or when applicants do not obfuscate. If uncertain examiners are lenient and grant patents with certainty, then increasing resources can backfire. Thus, our results suggest that the patent office ought to consider an art unit's leniency when allocating examination resources.

We present extensions of our baseline model that add more nuances to the framework. First, we allow applicants to obfuscate and *clarify* their ideas. In this setting, applicants with

[2] See "Updates to patent examination time, application routing, and examiner performance appraisal" available at (visited, December 2022): https://bit.ly/2WNGa1T

bad ideas obfuscate and applicants with good ideas clarify, and their incentives are driven by how uncertain examiners treat applications. Clarification goes in the opposite direction of obfuscation: The more likely an uncertain examiner to grant a patent, the higher the incentive to obfuscate, and the lower the incentive to clarify. If uncertain examiners always grant patents, there is no need for clarification; if they always reject, applicants with good ideas have the strongest incentives to clarify. We characterize the equilibrium in this setting, and show that the examiner's payoff is still non-monotone in the level of resources.

Second, we augment our model by adding more types of ideas (a continuum of values). This allows us to analyze how a change in the distribution of ideas impacts obfuscation. For instance, is there more or less obfuscation in settings with more "breakthrough" ideas (i.e., high value with small probability)? We show that a first-order stochastic dominance change in the distribution of ideas increases obfuscation but it is beneficial for the examiner. Reshaping the right tail of the distribution of ideas may increase or decrease obfuscation, depending on whether it increases or decreases the expected the value of sufficiently good ideas.

We then return to the baseline model and endogenize the proportion of good ideas by allowing applicants to invest in "knowledge," which increases the probability of coming up with a good idea. How much investment will there be? Too much investment make good ideas very abundant, so an uncertain examiner gives the benefit of the doubt (i.e., the examiner becomes lenient). Since acquiring knowledge is costly, an applicant may consider the alternative of not investing, and obfuscating whenever she draws a bad idea, given that uncertain examiners will grant a patent. This force dampens the incentive to invest in knowledge, making good ideas more scarce, reducing the leniency of examiners. In equilibrium, the proportion of good ideas balances out this trade-off.

Even after endogenizing the proportion of good ideas the relationship between examination resources and the examiner's payoff is non-monotonic. That is, increasing examination resources can still backfire. However, now there is an additional effect coming from the impact of more resources in the acquisition of knowledge, which determines the equilibrium proportion of good ideas. We show that more resources increases the equilibrium proportion of good ideas, since examiners learn the value of good ideas more frequently. A larger fraction of good ideas always benefit the examiner, so more resources has an additional benefit when the proportion of good ideas is determined endogenously.

Motivated by recent empirical evidence (Jones, 2009; Bloom et al., 2020), we study strategic obfuscation when good ideas are harder to find through three different channels. First, a

reduction in "low-hanging fruit," which makes it harder to find good ideas without investment. Second, less productive knowledge, i.e., good ideas are hard to come by for those who invest. Third, a more costly path to reach the frontier of knowledge (the "burden of knowledge," Jones, 2009). These channels have different implications on obfuscation, ex-post allocative efficiency, and ex-ante investment. Obfuscation decreases when ideas get harder to find except when the channel by which ideas get harder to find is a reduction of low-hanging fruit.

A reduction of low-hanging fruit has an ambiguous effect on obfuscation. A reduction of low-hanging fruit directly motivates agents to acquire knowledge. On the other hand, good ideas mostly come from agents who invest, since those who do not are less likely to find good ideas. These two effects combined have an ambiguous impact on the equilibrium proportion of good ideas and, therefore, on obfuscation. By contrast, obfuscation decreases both with lower productivity of knowledge or when investing is costlier because good ideas become scarcer. In the former case, investing becomes less attractive, so fewer agents invest, and those who do find good ideas less frequently. In the latter case, fewer agents specialize because it is more costly to do so.

Evidence of Obfuscation. Obfuscation has been empirically documented in various settings such as patent examination, loan approval, accounting, academia, and grant funding. Freilich (2017) documents the proliferation of language that "clutter" patent claims, and Kuhn et al. (2020) document an increase in (possibly irrelevant) prior-art citations in patent applications. Larrimore et al. (2011) find that verbosity and the expression of certainty increase loan funding success. The International Accounting Standards Board asserted that "the increasing size of annual reports had become a concern to the financial reporting community, with many key messages about a company's performance becoming drowned in detail."[3] Importantly, not every *type* of firm obfuscates. Li (2008) and Bushee et al. (2018) find that annual reports of firms with *lower* earnings are harder to read, i.e., poorly performing firms may want to bury the facts hoping that investors give them the benefit of the doubt.

Obfuscation in academic publications takes many forms (Nosek et al., 2012). There is widespread evidence of selective reporting and linguistic obfuscation (Head et al., 2015). Vinkers et al. (2015) find an 880 percent increase in the use of positive words in abstracts over 40 years, and Markowitz and Hancock (2016) find that fraudulent papers are written with significantly higher levels of linguistic obfuscation.

[3] https://bit.ly/3yJtQgp. Also, Drake et al. (2019) find that many professional financial-statement users believe disclosure overload is a problem.

Obfuscation is also prevalent in scientific grants, which are crucial to promoting private-sector innovation (Li, 2017; Howell, 2017; Azoulay and Li, 2020).[4] Markowitz (2019) finds that longer abstracts, written with more verbal certainty and fewer common words, receive more NSF funding. Importantly, obfuscation is more prevalent for low-quality proposals. Chubb and Watermeyer (2017) find that grant applicants feel the need to exaggerate and embellish the impact of their work, and more so when the impact of a research idea is not obvious.

Related Literature. Our framework relates to the literature on strategic communication. Crawford and Sobel (1982) study communication with "soft" information, whereas Milgrom and Roberts (1986) consider hard information. In other settings, the "softness" of information is endogenous: Legros and Newman (2002) study interference in court proceedings, where an agent can prevent the judge from receiving a message from her opponent; Dewatripont and Tirole (2005) study implications of different modes of communication. Similar to these articles, our setting considers endogenous softness of information. In contrast, we characterize the equilibrium, derive a number of comparative statics results, and study the implications of both an endogenously-determined distribution of private information and competition among senders.

Stahl and Strausz (2017) study the impact of having the seller or the buyer initiate third-party certification. They show that sellers use certification as a signaling device, which increases market transparency relative to certification from buyers who use it as an inspection device. In our setting there is no selection because all proposals must go through the reviewer. Perez-Richet and Prady (2012) present a related model where the receiver observes the sender's complexity choice but they do not study ex ante investment in knowledge nor competition among senders.

Our paper relates to the literature on project selection. Bar and Gordon (2014) uses a mechanism design approach to determine which projects to fund. In Che et al. (2013), the applicant must convince the sponsor that her project is worth it, without knowing the sponsor's outside option. Boleslavsky et al. (2019) explores funding decisions with information acquisition, although their focus is not on strategic obfuscation. In Boleslavsky and Cotton (2018), there are competing proposals and applicants provide evidence to convince the sponsor their projects deserve funding. In these papers, in contrast to our setting, senders choose signal structures (as in Bayesian persuasion), and find that competition can lead to more obfuscation.

[4] More than 80 percent of the National Institute of Health's $40 billion annual budget is allocated through competitive grants. See https://www.nih.gov/about-nih/what-we-do/budget.

Kartik et al. (2007) and Kartik (2009) study the incentives to costly misrepresent information. In contrast, in our setting the sender *cannot* mimic any type she wants but instead can prevent information transmission by obfuscating the reviewer. In Frankel and Kartik (2019) the sender's signaling cost is multidimensional, and the receiver cannot learn the sender's type—she only makes an inference based on a signal—whereas in our model she can learn it. So, in our model, the sender's effort *does not* signal a better type. Instead, it disrupts communication.

The industrial organization literature has studied firms that strategically obfuscate to hinder price discovery by consumers (Ellison and Wolitzky, 2012; Gu and Wenzel, 2014), or to exploit their limited attention (Persson, 2018). In our model, grant applicants (rather than firms) strategically obfuscate a reviewer who has limited resources (rather than consumers) .

A recent literature has measured the extent of obfuscation in controlled experiments. In Jin et al. (2021), senders cannot misrepresent information, but they can complicate a message. They find that complexity in disclosure is driven by the inability of receivers to fully internalize complex information. Similarly, de Clippel and Rozen (2020) also find that senders strategically obfuscate their messages.

2 Model: Strategic Obfuscation and Screening

There are two players: an examiner and a patent applicant. The examiner's goal is to grant patents only to novel, useful, and non-obvious ideas (henceforth, "good ideas"). The applicant is privately informed about the type v of the idea, which can be good ($v = V$) or bad ($v = 0$). The applicant communicates her idea in a patent application. The examiner reviews the application and either issues a patent or rejects it. An applicant of type v gets a payoff of $g(v) > 0$ from a patent and zero from rejection. The examiner derives a payoff of V from granting a patent to a good idea and incurs a cost c from granting a patent to a bad idea. A priori, the probability that an idea is good is $p = \Pr(v = V)$, which is commonly known.

The applicant chooses an *obfuscation* level $e \geq 0$ at cost $C(e) = e$, which makes it more difficult for the examiner to learn v. The examiner devotes a fixed amount of resources, $\sigma \in [0, 1]$, to examination. With probability $1 - \pi(\sigma, e)$ the examiner learns v. In that case, the examiner rejects a bad idea and grants a patent to a good one. With probability $\pi(\sigma, e)$ the examiner does not find conclusive evidence and is *uncertain* about the idea's value after examination.

An uncertain examiner holds a belief Q that the idea is good. Based on this belief, she chooses to grant a patent or reject the application in order to maximize her expected payoff. When indifferent, she grants the patent with probability α. Figure 1 summarizes the timing of the events in our model.

Ideas	Applicant	Examiner's signal	Examiner's decision
Value $v \in \{0, V\}$ to the sponsor. $\Pr(v = V) = p$.	Privately observes v. Chooses *obfuscation* effort $e(v)$.	Supplies examination time σ. After examination either learns $v \in \{0, V\}$ or is uncertain.	Grant a patent if finds $v = V$. Reject if finds $v = 0$. If uncertain, decide based on Q, and if indifferent, grant a patent with probability α.

Figure 1: Timeline of the game.

Assumption 1. $\frac{\partial \pi(\sigma,e)}{\partial \sigma} \leq 0 \leq \frac{\partial \pi(\sigma,e)}{\partial e}$, *and* $\pi(1,0) = 0 < \pi(\sigma,0)$ *for* $\sigma < 1$. *Also,* $\frac{\partial^2 \pi(\sigma,e)}{\partial e \partial e} \leq 0$ *and* $\frac{\partial^2 \pi(\sigma,e)}{\partial \sigma \partial e} > 0$.

Obfuscation impairs the examiner's ability to find hard evidence, so $\pi(\sigma, e)$ increases in e. However, conditional on an obfuscation level, an examiner with more resources is less likely to be uncertain, so $\pi(\sigma, e)$ decreases in σ. The condition $\pi(1, 0) = 0 < \pi(\sigma, 0)$ for $\sigma < 1$ means that even if there is no obfuscation an imperfect examiner ($\sigma < 1$) is uncertain with positive probability, whereas a "perfect" examiner, indexed by $\sigma = 1$, always finds the idea's value. Lastly, we assume non-increasing marginal returns to obfuscation (weak concavity in e), and the supermodularity condition reflects that the marginal return to obfuscation is larger when examiners are more resourceful. Assumption 1 is satisfied, for instance, when $\pi(\sigma, e) = 1 - a(\sigma)q(e)$, with $a(\sigma)q(e) \in [0, 1]$, $q'(\cdot) < 0 < a'(\cdot)$, and $q''(e) \geq 0$.

Discussion of the Framework and Extensions. We capture a conflict of interest between a decision maker and a privately informed sender in a setting where communication is probabilistic, as in Dewatripont and Tirole (2005). Importantly, the sender cannot misrepresent her information; she can exert effort to disrupt communication. This framework captures that the examiner can make both type-1 and type-2 errors.

Our baseline framework makes a number of assumptions that simplify the analysis. However, we can extend it in multiple directions and show that these more complex models qualitatively preserve our main results. We discuss these extensions below.

Clarification and Obfuscation. Applicants can only choose to obfuscate in our baseline model. In practice, some applicants can obfuscate and others can clarify. In Section 4.1, we extend

our framework to allow for both obfuscation and clarification effort.

Richer Type Space. Our baseline model assumes only two types: applicants with good or with bad ideas. In practice, ideas are heterogeneous, and their quality belongs to a spectrum. In Section 4.2, we study an extension with a continuum of types. This extension allows us to ask questions related to changes in the distribution of ideas. For instance, is there more or less obfuscation when "breakthrough" ideas (those with extremely high value and low probability) are more likely?

Examiner effort. We assume that the examiner's resources are fixed: $\sigma \in [0, 1]$ represents resources (e.g., hours of examination). This assumption is justified by the weak incentives faced by patent examiners whose salaries depend on meeting quotas, and they are not liable for granting patents that are ex post invalid (Langinier and Marcoul, 2020). Time-constrained examiners face convex costs from exerting effort on any given application. Fixed resources implies extreme convexity: Using up to σ resources is costless, but beyond σ is prohibitively costly. In Online Appendix A, we show that the equilibrium taxonomy is similar when the examiner faces explicit performance incentives, and strategically chooses examination resources at a quadratic cost.

Richer Signal Space. Our main analysis follows Dewatripont and Tirole (2005) in that the examiner either finds hard evidence or not. More generally, examination might produce different degrees of signals correlated with the value of the idea. In Online Appendix B, we show that our main results hold in that richer environment.

Ex-ante Commitment. Since our game is dynamic, a natural question is whether there is value to commit to a strategy. In Online Appendix C, we study a framework where the examiner commits ex-ante to granting patents with a fixed probability whenever uncertain. We show that commitment weakly benefits the examiner, although in many cases it does not produce any benefit.

Submodularity. Our baseline model assumes larger marginal return to obfuscation when applicants face more resourceful examiners. In Online Appendix D, we discuss how results change if we reverse this assumption, i.e., $\pi(\cdot)$ was submodular rather than supermodular.

3 Equilibrium Analysis: When will there be obfuscation?

We use perfect-Bayesian equilibrium (PBE) as the solution concept.

Examiner. After examining the application, the examiner either learns the idea's value or remains uncertain about it. If she learns v the decision is trivial: she grants a patent if $v = V$, and rejects the application when $v = 0$. If she does not learn it and is uncertain, she believes that the idea is good with probability Q. Based on this belief, she grants a patent with probability $\alpha \in [0, 1]$ where

$$\alpha \in \arg\max_{\alpha \in [0,1]} \alpha \cdot (QV + (1 - Q)(-c)).$$

In choosing α, the examiner trades off the expected benefit of granting a patent to a good idea versus the expected cost of granting a patent to a bad idea. In particular, note that adopting the strategy of never granting a patent when uncertain might be suboptimal.

Applicant's Optimal Obfuscation Effort. An applicant with a good idea never obfuscates because she wants the examiner to learn that $v = V$, and obfuscation can only hamper learning. However, an applicant with a bad idea considers obfuscation because making the examiner uncertain is weakly better than letting her learn that the idea is bad. Let $g \equiv g(0)$. The optimal level of obfuscation for an applicant with a bad idea is

$$e(\alpha, \sigma) \in \arg\max_{e \geq 0} \alpha \cdot g \cdot \pi(\sigma, e) - e. \tag{1}$$

With probability $\pi(\sigma, e)$, the examiner is uncertain about the value of the idea after examining the application and grants a patent with probability α, in which case the applicant with a bad idea receives a payoff of g. The applicant with a bad idea does not obfuscate when the marginal return to obfuscation is low enough. Specifically, when $g\alpha\pi_e(\sigma, e) < 1$ for $e > 0$. Otherwise, the applicant's obfuscation effort is positive and increases with α and g. By Assumption 1, the marginal return from obfuscation is higher when applicants face more resourceful reviewers, which implies (by Topkis' monotonicity theorem) that $e(\alpha, \sigma)$ is increasing in σ.

A complete description of the equilibrium includes the optimal obfuscation effort, e^*, the sponsor's belief, Q^*, and the probability of funding uncertified proposals, α^*.

Definition 1. *An equilibrium is characterized by* (e^*, α^*, Q^*) *that satisfy:*

1. *Applicants with bad ideas obfuscate optimally:* $e^* = e(\alpha^*, \sigma)$.

2. *The posterior belief Q^* is consistent with Bayes' rule:*

$$Q^* = Pr(v = V|uncertain) = \frac{\pi(\sigma, 0)p}{\pi(\sigma, 0)p + \pi(\sigma, e^*)(1-p)}. \tag{2}$$

3. *The probability of granting a patent is consistent with the sponsor's posterior belief:*

$$\alpha^* \in \arg\max_{\alpha \in [0,1]} \alpha \cdot (Q^* V + (1 - Q^*)(-c)).$$

Granting a patent creates an expected benefit and cost for an uncertain examiner. Out of the applications that leave the examiner uncertain, a proportion $p\pi(\sigma, 0)$ are good ideas for which a patent that creates a value V, and a proportion $(1-p)\pi(\sigma, e^*)$ are bad ideas for which a patent creates a loss of c. The benefit is larger than the cost if and only if

$$p\pi(\sigma, 0)V > (1-p)\pi(\sigma, e^*)c, \tag{3}$$

which is equivalent to $Q^*(V + c) > c$. This tradeoff determines the type of equilibrium:

Proposition 1. *Let Q^* be the examiner's posterior belief. In equilibrium we have:*

1. *$Q^*(V + c) > c$, an uncertain examiner grants a patent ($\alpha^* = 1$); applicants with bad ideas obfuscate by $e^* = e(1, \sigma)$.*

2. *$Q^*(V + c) < c$, an uncertain examiner rejects the application ($\alpha^* = 0$); applicants with bad ideas do not obfuscate, $e^* = 0$.*

3. *If $Q^*(V + c) = c$, an uncertain examiner is indifferent between granting a patent or rejecting the application. The examiner grants a patent with probability $\alpha^* \in [0, 1]$ and applicants with bad ideas obfuscate by $e^* = e(\alpha^*, \sigma)$. The probability α^* is pinned down by Bayes rule and satisfies*

$$\pi(\sigma, 0)pV = \pi(\sigma, e(\alpha^*, \sigma))(1-p)c. \tag{4}$$

Using Proposition 1 and equation (17) we can characterize the PBE for any set of exogenous parameters. For the sake of exposition, we relegate the details of this characterization to the proof of Proposition 1 and use Figure 2 to illustrate three qualitatively different regions in the exogenous parameters space, (p, σ). First, if $p(V + c) \leq c$ or, equivalently, $pV < (1-p)c$, the examiner would never grant a patent prior to examination. If uncertain, the examiner revises her belief downwards, so never grants a patent, and there is no point in obfuscating.

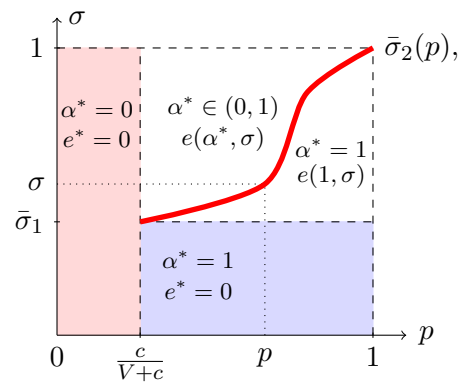

Figure 2: There is no obfuscation effort in the shaded region. For $p(V+c) > c$, the (red) curve shows $\bar{\sigma}_2(p)$. The threshold $\bar{\sigma}_2(p)$ is formally defined by setting $\alpha^* = 1$ in (17) and solving for σ. When $\sigma \in (\bar{\sigma}_2(p), 1]$ there is randomization, $\alpha^* \in (0,1)$ is pinned down in (17). When $\sigma \in (\bar{\sigma}_1, \bar{\sigma}_2(p)]$, all uncertified projects are funded, $\alpha^* = 1$. Equilibrium obfuscation is $e^*(\alpha^*, \sigma)$ pinned down in (1).

This corresponds to situations where good ideas are scarce or the value of patenting a good idea is low relative to the cost of patenting a bad one.

There is also no obfuscation when σ is below a threshold $\bar{\sigma}_1$.[5] Below this threshold, the examiner's resources are so low that she is uncertain after reviewing most of the applications. Since obfuscation is costly, the marginal benefit of obfuscation is lower than its marginal cost.

Second, if $p(V+c) > c$, there exists a unique threshold $\bar{\sigma}_2(p) \in (\bar{\sigma}_1, 1)$—the red curve in Figure 2—such that inequality (3) holds when $e^* = e(1, \sigma)$ and $\sigma \in (\bar{\sigma}_1, \bar{\sigma}_2]$. The threshold $\bar{\sigma}_2$ is the maximum level of resources for which applicants exert the maximum obfuscation level, $e^* = e(1, \sigma)$, and the examiner grants a patent when she is uncertain. In Figure 2, $p(V+c) > c$ and $\sigma \in (\bar{\sigma}_1, \bar{\sigma}_2]$ correspond to the region below the line, where $\alpha^* = 1$ and $e^* = e(1, \sigma)$. This region capture settings where ex-ante ideas are valuable, $pV > (1-p)c$, and examiners resources are limited and do not change the examiner's belief too much, even when applicants exert the maximum obfuscation level, $e(1, \sigma)$. Thus, uncertain examiners always grant patents.

Third, if $p(V+c) > c$ and $\sigma > \bar{\sigma}_2(p)$, an uncertain examiner grants a patent with probability α^* strictly less than one, pinned down by (17). If an uncertain examiner granted a patent for sure, applicants with bad ideas would obfuscate too much, and the examiner would be better off rejecting. Instead, applicants obfuscate up to the point at which the examiner is indifferent

[5] Formally, $\bar{\sigma}_1$ is defined as the maximum value of σ such that the solution to (1) is $e^* = 0$ for all $\sigma < \bar{\sigma}_1$. The value of $\bar{\sigma}_1$ depends on the shape of $\pi(\cdot)$, and on the paramters p, V and c.

between granting a patent or rejecting the application. When indifferent, the examiner picks $\alpha^* \in [0,1]$ to satisfy (17). In Figure 2, these cases correspond to region where $\alpha^* \in (0,1)$ and $e^* = e(\alpha^*, \sigma)$.

We next illustrate α^* and e^* as functions of σ, fixing all the other parameters such that $p(V+c) > c$.[6] In Figure 3a, for $\sigma < \bar{\sigma}_1$, there is no obfuscation and uncertain examiners reject. For $\sigma \in (\bar{\sigma}_1, \bar{\sigma}_2)$, although there is obfuscation, $e^* = e(1,\sigma) > 0$, an uncertain examiner's benefit of granting a patent outweighs the cost of rejecting, so $\alpha^* = 1$. When $\sigma > \bar{\sigma}_2$, the obfuscation effort is $e^* = e(\alpha^*, \sigma)$ and an uncertain examiner grants a patent with probability $\alpha^* < 1$.

Lastly, Figure 3b illustrates α^* and e^* as functions of p, fixing all the other parameters such that $\sigma > \bar{\sigma}_1$.[7] When $p(V+c) \leq c$, $\alpha^* = 0$ and $e^* = 0$. For $p > \bar{p}(\sigma)$ ($\bar{p}(\cdot)$ is the inverse of $\bar{\sigma}_2(\cdot)$), inequality (3) holds so uncertain examiners grant a patent ($\alpha^* = 1$). For $p \in \left(\frac{c}{V+c}, \bar{p}(\sigma)\right)$, $\alpha^* \in (0,1)$ solves (17), i.e. it equates the benefit and the cost of granting a patent. Obfuscation increases in p until $\bar{p}(\sigma)$, reaching its the maximum level, $e^* = e(1,\sigma)$, for $p > \bar{p}(\sigma)$.

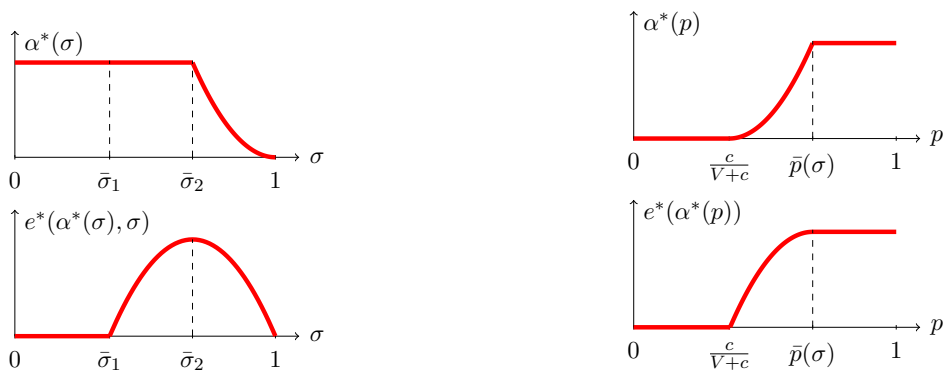

(a) α^* and e^* as a function of σ, when $pV > c$.

(b) α^* and e^* as a function of p, when $\sigma > \bar{\sigma}_1$

Figure 3: Equilibrium comparative statics, for α^* and e^*, in σ and p.

3.1 Empirical Connection

Our model predicts that applicants with bad ideas obfuscate. Is there evidence of strategic obfuscation? Ashtor (2022) shows that patent applicants strategically obfuscate.[8] Using a

[6] If $p(V+c) \leq c$, we have $e^* = 0$ and $\alpha^* = 0$.

[7] If $\sigma \leq \bar{\sigma}_1$, then $e^* = 0$. $\alpha^* = 0$ for $p(V+c) \leq c$ and $\alpha^* = 1$ for $p(V+c) > c$.

[8] They define a measure of clarity $y \in [0,1]$. A measure of obfuscation is $\tilde{y} = 1 - y$.

2014 Supreme Court decision that increased the cost of obfuscation, they find that applicants decrease their obfuscation after the policy change. They also show that higher-value patents tend to be less obfuscated, which is consistent with our model. Following Ashtor (2022), we use linguistic features of patent applications to measure obfuscation.

We collect data on all patent applications filed in the U.S. patent office between 2015 and 2017 to empirically test the predictions of our model that applicants obfuscate more when facing more lenient examiners (higher α^*). Examiners can be more lenient in technology areas where they have lower resources (see Figure 3a) or where the prior that an idea is good is higher (see Figure 3b).

First, we use all the U.S. patent applications filed in 2015 to construct a proxy measure of art unit "leniency," which we define as the fraction of applications in the art unit that eventually get a patent. Figure 4 shows a large variation of leniency across art units. For instance, art units in the technology center "Digital and Optical Communications" are at the right of the distribution (e.g., art unit 2633), patenting over 90 percent of all applications. At the other extreme, art units associated to "Financial and Business practices" (e.g., art unit 3685) patent less than 20 percent of all applications.

According to our model, applicants facing more lenient art units (e.g., higher α^*) will obfuscate more. To empirically test this prediction, we construct obfuscation measures for patent applications filed in 2016 and 2017, including measures of obfuscation in the application claims (indefininess, inclusiveness, number of words) and in the description of the invention (concreteness, inclusiveness, number of words). Since it is computationally expensive to implement all the linguistic features in Ashtor (2022), for the purpose of this paper, we implement only a few of them. We construct a measure of "indefinitness." Examiners must decide whether a person "skilled in the art" would understand what the applicant claims is the invention's contribution when the claim is read in light of the specification. If it is ambiguous, the claim is "indefinite," and certain keywords make claims more prone to indefiniteness: means for, step for, mechanism for, module for, device for, unit for, component for, element for, member for, apparatus for, machine for, system for. We measure "indefiniteness" as the count of how many of these keywords appear in the claims of an applications.

We also construct measures that capture the cognitive burden the examiner is subject to when reading an application. We construct measures of "inclusiveness" both for claims and the description of the text. Higher inclusiveness is associated with a higher cognitive burden (see, e.g., Iliev and Axelrod, 2017). Additionally, we construct a measure of "concreteness"

for the description of the text. Lower concreteness makes the text more ambiguous, harder to read. We also count the number of words both in the description of the invention and the claims.

Lastly, we compute a measure of overlap between the description of the invention and the claims. Narrow overlap makes it more difficult for the examiner to understand what the applicant claims to be the contribution of the invention.

Using these measures of obfuscation, we test whether the "leniency" at each art unit correlates (and if so, in which direction) with these measures of obfuscation. We estimate the following linear regression model

$$Y_{it} = \alpha + \beta \cdot \text{Leniency}(2015)_{t(i)} + \varepsilon_{it},$$

where Y_{it} is a measure of obfuscation of interest for an application i (filed in 2016 or 2017) that belongs to art unit $t(i)$, and $\text{Leniency}(2015)_{t(i)}$ is the patenting rate of art unit $t(i)$. Table 1 reports our results, which are consistent with our model's prediction that higher α induces more obfuscation.

Columns 1 to 3 suggest that obfuscation at the claim level positively correlate with leniency: Applications in art units with higher prior patenting rates are more indefinite, inclusive, and have more words. Columns 4 to 6 suggest that obfuscation in the description of the invention also positively correlate with the prior probability of acceptance: Applications in art units with a higher prior patenting rates are less concrete, more inclusive, and have less words. Note that more words in the claims and less words in the description are both consistent with a higher burden to understand whether the invention is a good idea or not. Lastly, Column 7, shows less overlap between the description of the invention and the text in the claims in more lenient art units, suggesting that it is harder to assess the contribution of the invention.

These results together suggest that applicants with inventions in more lenient art units (e.g., higher α in our model) are more likely to obfuscate.

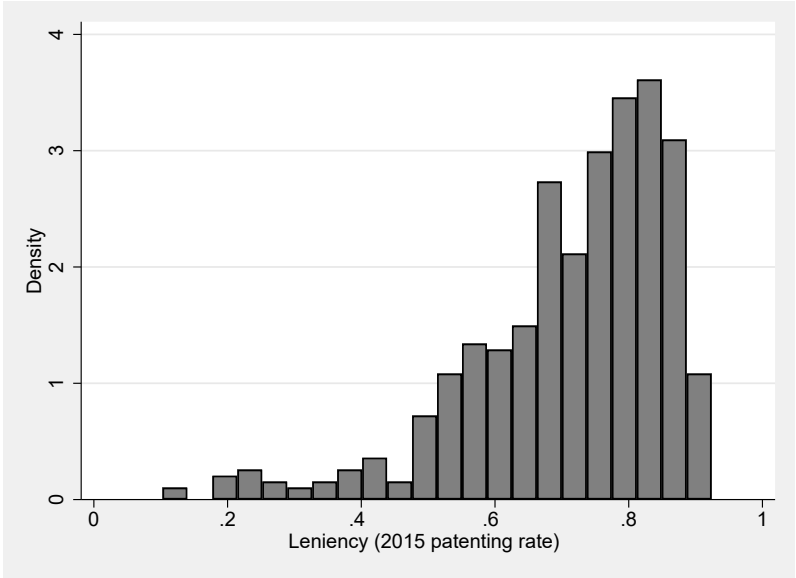

Figure 4: Rate of patent granting across art units.

Table 1: Regression Results

	Claims			Description			Overlap
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Indefiniteness	Inclusiveness	Word Count	Concreteness	Inclusiveness	Word Count	Jaccard
Leniency (2015)	0.261***	0.136***	61.47***	-0.057***	0.160***	-142.6***	-0.00798***
	(0.024)	(0.004)	(7.168)	(0.000)	(0.003)	(12.030)	(0.000)
Mean	0.73	15.10	1007.70	2.48	15.23	4030.40	0.028
Obs	658,342	658,334	658,342	658,340	658,332	658342	658,342

Notes: Standard errors in parentheses. $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$

3.2 More Resources for Patent Examiners may Backfire

Concerns about the quality of the patent examination—see, e.g., Frakes and Wasserman (2017)—have triggered a series of recent reforms in the United States Patent and Trademark Office, which provided more resources to examiners (see footnote 2). Our analysis shows that endowing examiners with more resources creates countervailing effects. All-else equal examiners are less likely to be uncertain after examination. However, applicants obfuscate more when they face more resourceful examiners.

The examiner's equilibrium payoff is

$$u_E = p(1 - \pi(\sigma, 0))V + \alpha^*[p\pi(\sigma, 0)V - (1 - p)\pi(\sigma, e^*)c]. \tag{5}$$

A fraction $p(1-\pi(\sigma,0))$ of the applications are good ideas that receive a patent because the examiner learns that $v = V$. A fraction $(1-p)(1-\pi(\sigma,e^*))$ are bad ideas that are rejected because the examiner learns that $v = 0$. The rest of the applications leave the examiner uncertain. If an uncertain examiner grants a patent with probability α^*, she receives an expected payoff of $p\pi(\sigma,0)V - (1-p)\pi(\sigma,e^*)c$.

Whenever $\alpha^* \in (0,1)$ the examiner must be indifferent between granting a patent and rejecting the application, so $p\pi(\sigma,0)V = (1-p)\pi(\sigma,e^*)c$. Thus, we can write the examiner's payoff as:

$$u_E = \begin{cases} p(1-\pi(\sigma,0))V & \text{, if } \alpha^* < 1, \\ pV - (1-p)\pi(\sigma,e^*(1,\sigma))c & \text{, if } \alpha^* = 1. \end{cases}$$

The impact of increasing resources on the examiner's payoff can be divided into three cases. First, when an uncertain examiner never grants a patent ($\alpha^* = 0$), there is no obfuscation. Hence, a marginal increase in σ always benefits the examiner because only the first order effect of finding good ideas more frequently matters. This occurs, for instance, when good ideas are scarce.

Second, when an uncertain examiner is indifferent between granting the patent or rejecting the application, obfuscation is irrelevant since an uncertain examiner's expected payoff is zero. Hence, here again a marginal increase in σ always benefits the examiner because only the first order effect of finding good ideas more frequently matters.

Lastly, when an uncertain examiner grants a patent for sure ($\alpha^* = 1$), all good ideas are patented, so more resources only matter to the extent to screening our bad ideas. A marginal increase in σ creates two opposing forces: the examiner is less likely to be uncertain but applicants obfuscate more since the marginal return to obfuscation increases in σ.

Proposition 2. *Increasing resources for a lenient examiner can backfire: The examiner's payoff is* decreasing *in σ if and only if $p(V+c) > c$, $\sigma \in (\bar{\sigma}_1, \bar{\sigma}_2(p))$ and*

$$\frac{\partial \pi(\sigma,e^*)}{\partial \sigma} + \frac{\partial \pi(\sigma,e^*)}{\partial e}\frac{\partial e^*}{\partial \sigma} > 0. \tag{6}$$

Inequality (6) shows the net effect of two countervailing effects. First, examination leaves the examiner uncertain less often when σ is higher, which is captured by the negative term $\frac{\partial \pi(\sigma,e^*)}{\partial \sigma}$. Second, facing more resourceful examiners, applicants obfuscate more, which is captured by the positive term $\frac{\partial e^*}{\partial \sigma}$, and more obfuscation leaves the examiner uncertain more often, which is captured by the positive term $\frac{\partial \pi(\sigma,e^*)}{\partial e}$. More resources make the examiner worse off when the latter effect dominates the former.

We can write (6) in terms of elasticities to ease the interpretation

$$\varepsilon_{\pi,\sigma} + \varepsilon_{\pi,e}\varepsilon_{e,\sigma} > 0.$$

A 1% increase in σ decreases the probability that the examiner is uncertain by $\varepsilon_{\pi,\sigma}$ percent, and increases this probability (indirectly, through obfuscation) by $\varepsilon_{\pi,e}\varepsilon_{e,\sigma}$ percent. Increasing resources leaves the examiner uncertain more often when the net effect is positive. Whether (6) holds depends on the examination technology, $\pi(\cdot)$ and other model parameters.

Figure 5 illustrates the non-monotonicity of the examiner's payoff as a function of σ when $p(V + c) > c$ and (6) holds. Figure 5a identifies two thresholds, $\sigma_{\min}$ and $\sigma_{\max}$, such that the examiner prefers $\sigma_{\min}$ to any $\sigma \in (\sigma_{\min}, \sigma_{\max})$, i.e., the examiner can be better off with *less* resources. Increasing σ beyond $\sigma_{\max}$ always improves the examiner's payoff. Of course, if unlimited resources were available at no cost, the examiner would choose $\sigma = 1$. In practice, it could be too costly (even impossible) to procure examination resources beyond $\sigma_{\max}$. Suppose that there is a strictly increasing and convex cost $\kappa(\cdot)$ of procuring more resources. Then, the optimal level of resources, σ^*, maximizes $u_E(\sigma) - \kappa(\sigma)$.

Figure 5b the examiner's payoff in a simulation where $\pi(\sigma, e) = 1-\sigma q(e)$, where $q(e) = 1-\sqrt{2e}$ for $e \leq 0.5$ and zero otherwise; $p = 0.7$, $V = 2$, and $c = 1$. With these parameters, $\bar{\sigma}_1 = 0$ and $\bar{\sigma}_1 = 0.817$. Moreover, (6) holds, $\sigma_{\min} = 0.5$ and $\sigma_{\max} = 0.84$.

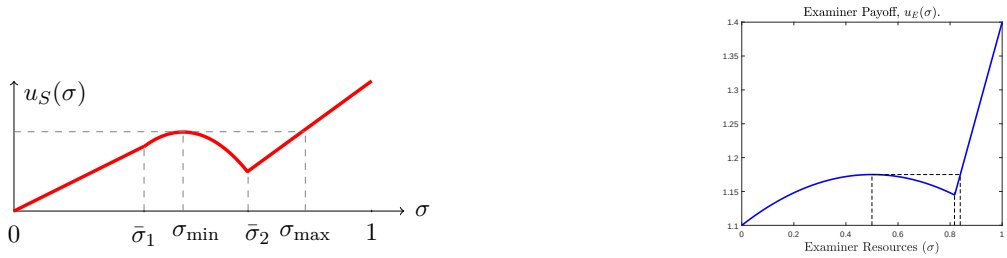

(a) Illustration of non-monotone payoff.

(b) Simulation Example.

Figure 5: Examiner's payoff as a function of σ, when $p(V + c) > c$ and (6) holds.

Empirical connection. Our results suggest that to empirically assess whether more examination resources are beneficial or detrimental, a researcher needs to estimate the resources-elasticity and the obfuscation-elasticity of examination. One way to proceed is to structurally estimate the function $\pi(\sigma, e)$ using application-level data. Once this function is known, we can assess whether (6) holds.

Alternatively, we need a "shock" to examination resources. A plausible shock is a recent

change in the USPTO, beginning on October 1st 2019, that increased resources for examiners. In our model, this change maps directly into an increase in σ. The policy roll-out was staggered, implemented in a few technology units first, and later, in 2020, in all of them.[9] We could use this policy change to measure how applicants change the obfuscation in their patent applications, comparing technology units where the change was implemented with those where the change was not implemented using, for instance, a difference-in-difference research design

$$Y_{iat} = \alpha + \beta_0 1(a \text{ treated}) + \beta_1 1(\text{after}) + +\beta_2 1(a \text{ treated}) 1(\text{after}) + \varepsilon_{ita},$$

where Y_{ita} is a measure of obfuscation for application i, which is an invention that belongs to art unit a and was filed in period t. $1(a \text{ treated})$ is an indicator that takes the value 1 if art unit a received more resources and 0 otherwise. $1(\text{after})$ is indicator that takes the value 1 if application i was filed after examiners received more resources and 0 otherwise. The coefficient of interest if β_2. If this coefficient is positive and significant, art units subject to more resources receive more obfuscated applications. Since this change was recent and there are delays in the examination process, we leave this exercise for future work.

4 Extensions

4.1 Clarification and Obfuscation

We consider a setting that incorporates both clarification and obfuscation. Here, there are non-trivial decisions by both types of applicants. We can easily extend our results to this setting, although it requires additional assumptions and more notation.

Let $\pi(\sigma, e_O, e_C)$ be the probability that an examiner is uncertain about the quality of an idea after examination when the applicant exerted obfuscation effort e_O and clarification effort e_C. We assume that this function is increasing in e_O and decreasing in e_C. Naturally, applicants with bad ideas never clarify, but might obfuscate, and applicants with good ideas might clarify, but never obfuscate.

A PBE equilibrium requires consistency between four variables: Q^*, α^*, e_O^*, and e_C^*.

[9] See pp. 68 in the 2021 USPTO's "Performance and Accountability Report," (visited, October 2022): https://www.uspto.gov/sites/default/files/documents/USPTOFY21PAR.pdf

1. Applicants with good ideas choose their clarification effort, e_C^*, to solve:

$$\max_{\alpha \in [0,1]} g[1 - \pi(\sigma, 0, e_C) + \alpha\pi(\sigma, 0, e_C)] - e_C$$

2. Applicants with bad ideas choose their obfuscation effort, e_O^*, to solve:

$$\max_{\alpha \in [0,1]} g\alpha\pi(\sigma, e_O, 0) - e_O$$

3. Given a posterior belief Q^*, the examiner grants a patent with probability α^* to solve:

$$\max_{\alpha \in [0,1]} \alpha(Q^*V - (1 - Q^*)c)$$

4. The examiner's posterior belief must be consistent with Bayes rule, given e_C^* and e_O^*:

$$Q^* = \frac{\pi(\sigma, 0, e_C^*)p}{\pi(\sigma, 0, e_C^*)p + \pi(\sigma, e_O^*, 0)(1 - p)}$$

We have three cases: $\alpha^* = 1$ is an equilibrium if $Q^*(V + c) > c$; $\alpha^* = 0$ is an equilibrium if $Q^*(V + c) < c$ and $\alpha^* \in (0, 1)$ is an equilibrium if $Q^*(V + c) = c$. Let

$$g(\alpha) \equiv Q^*(V + c) - c = \pi(\sigma, 0, e_C^*(\alpha, \sigma))pV - \pi(\sigma, e_O^*(\alpha, \sigma), 0)(1 - p)c.$$

Using this function, we have that $g(1) \geq 0$ implies $\alpha^* = 1$ is an equilibrium; $g(0) \leq 0$ implies $\alpha^* = 0$; and if it exists α^* such that $g(\alpha^*) = 0$, then α^* is an equilibrium. When $g(0) > 0$ and $g(1) < 0$, by continuity of $g(\cdot)$ an equilibrium exist.

In contrast to the baseline setting, where equilibrium exists and is unique, here we could potentially have multiple equilibria. To guarantee that a unique equilibrium exists we need to impose additional conditions. For example, uniqueness is guaranteed if $g(\cdot)$ satisfies the single crossing condition.

The examiner's payoff is

$$(1 - \pi(\sigma, 0, e_C^*)))pV + \alpha^*[\pi(\sigma, 0, e_C^*)pV - \pi(\sigma, e_O^*, 0)(1 - p)c],$$

which can be written as

$$\begin{cases} (1 - \pi(\sigma, 0, e_C(\alpha^*, \sigma)))pV & \alpha^* < 1, \\ pV - \pi(\sigma, e_O(1, \sigma), 0)(1 - p)c & \alpha^* = 1. \end{cases}$$

when $\alpha^* = 1$, the analysis is analogous to the baseline case, since $e_C^* = 0$. Hence, the result that more resources can make the examiner worse off is preserved.

Clarification leads to more nuanced results when $\alpha^* < 1$. In the baseline model, the examiner's payoff increases in σ in the region where $\alpha^* < 1$. With both clarification and obfuscation, however, there an additional effect, since applicants with good ideas clarify. Consider a region of parameters where $\alpha^* < 1$. The examiner's payoff increases when σ increases when

$$\frac{\partial \pi(\sigma, 0, e_C(\alpha^*, \sigma)))}{\partial \sigma} + \frac{\partial \pi(\sigma, 0, e_C(\alpha^*, \sigma)))}{\partial e_C} \left[\frac{\partial e_C(\alpha^*, \sigma)}{\partial \alpha} \frac{\alpha^*}{\partial \sigma} + \frac{\partial e_C(\alpha^*, \sigma)}{\partial \sigma} \right] < 0$$

The first term is the direct effect of increasing σ, $\frac{\partial \pi(\sigma, 0, e_C(\alpha^*, \sigma)))}{\partial \sigma}$, which is negative: more resourceful examiners are more able to identify good ideas. The second term is the indirect effect. $\frac{\partial \pi(\sigma, 0, e_C(\alpha^*, \sigma)))}{\partial e_C}$ is negative since clarification reduces the chance of leaving the examiner uncertain. Thus, the sign of the total effect hinges on the sign of the term inside the bracket. Thus, in contrast to the baseline case, clarification and obfuscation combined could make a more resourceful examiner worse off when $\alpha^* < 1$.

4.2 Richer Type Space

In this section, we extend the model to allow for a continuum of idea values. Applications contain ideas of different values to the examiner, with $v \sim F[0, \infty)$. A patentable idea is one that is novel, useful, and non-obvious. We capture this patentability requirement with a threshold $\bar{v} > 0$, so the examiner's goal is to grant patents only to ideas of value $v > \bar{v}$. We make two assumptions here: (1) the examiner's cost of granting a patent for ideas of value $v \leq \bar{v}$ is constant and equal to c; (2) applicants with ideas of value $v < \bar{v}$ receive the same private value from a patent. The justification of these assumptions is that the value that an applicant gets from patenting and idea of value v comes from rent extraction when negotiating a license under the shadow of litigation. Litigation is costly and entails imperfect re-examination, i.e., it does not invalidate the patent of an idea $v < \bar{v}$ with certainty. The value and social cost of a patent for a worthless idea comes from the licensing negotiation, which we assume identical for all patent associated to ideas $v < \bar{v}$.

As in the baseline case, if the examination process reveals v, the examiner's decision is trivial. However, if the examiner is uncertain, she makes a decision based on the belief $Q(\cdot)$, which is a (posterior) probability distribution over $[0, \infty)$. Based on this belief, she chooses the probability of granting a patent, α, by solving

$$\max_{\alpha \in [0,1]} \alpha \cdot E_Q \left(v 1\{v > \bar{v}\} - c 1\{v \leq \bar{v}\} \right).$$

Applicants with ideas of value $v > \bar{v}$ do not obfuscate, while applicants with ideas of value

$v \leq \bar{v}$ might obfuscate. They solve:

$$\max_{e \geq 0} \alpha g \pi(\sigma, e) - e.$$

All applicants with ideas $v \leq \bar{v}$ exert the same obfuscation level, $e^*(\alpha, \sigma)$.

In a perfect-Bayesian equilibrium, the examiner's belief must be consistent with Bayes rule:

$$dQ(v) = \frac{\pi(\sigma, 0)1(v > \bar{v}) + \pi(\sigma, e)1(v < \bar{v})}{\pi(\sigma, 0)(1 - F(\bar{v})) + \pi(\sigma, e)F(\bar{v})} dF(v)$$

Then,

$$E_Q\left(v1\{v > \bar{v}\} - c1\{v \leq \bar{v}\}\right) = \frac{\pi(\sigma, 0)E[v|v > \bar{v}](1 - F(\bar{v})) - c\pi(\sigma, e)F(\bar{v})}{\pi(\sigma, 0)(1 - F(\bar{v})) + \pi(\sigma, e)F(\bar{v})},$$

so the equilibrium is:

1. $\alpha^* = 1$ if $\pi(\sigma, 0)E[v|v > \bar{v}](1 - F(\bar{v})) > c\pi(\sigma, e(1, \sigma))F(\bar{v})$.
2. $\alpha^* = 0$ if $\pi(\sigma, 0)E[v|v > \bar{v}](1 - F(\bar{v})) < c\pi(\sigma, e(0, \sigma))F(\bar{v})$.
3. $\alpha^* \in [0, 1]$, where α^* is pinned down by $\pi(\sigma, 0)E[v|v > \bar{v}](1 - F(\bar{v})) = c\pi(\sigma, e(\alpha^*, \sigma))F(\bar{v})$.

The examiner's equilibrium expected payoff is

$$(1 - F(\bar{v}))(1 - \pi(\sigma, 0))E[v|v > \bar{v}] + \alpha^*[\pi(\sigma, 0)E[v|v > \bar{v}](1 - F(\bar{v})) - c\pi(\sigma, e(\alpha^*, \sigma))F(\bar{v})]$$

The same argument in the baseline model applies to show that only the case $\alpha^* = 1$ matters when looking at the comparative static of σ. Moreover, the derivative with respect to σ is the same as in the baseline model. Therefore, giving more resources to examiners can also be detrimental in a setting with a continuum of types.

Changes in the Distribution of Values. Some ideas are "breakthroughs" defined as very high-value ideas that occur with very small probability. How does the existence of breakthrough ideas affect the incentive to obfuscate and to grant patents? We can answer this question by analyzing the impact of a change in F on obfuscation. Applicants with ideas $v > \bar{v}$ do not obfuscate. However, a change in the distribution of idea values affect an uncertain examiner's decision to grant a patent or reject an application, which in turn affects the applicants incentives to obfuscate.

The key to analyze the impact of a change in the distribution F on obfuscation is the term:

$$\mathrm{R}(F) = \frac{E[v|v > \bar{v}](1 - F(\bar{v}))}{F(\bar{v})}. \tag{7}$$

Proposition 3. *Consider two distributions, F and G, with $R(F) \geq R(G)$. Then, there is weakly more obfuscation under F than under G.*

Thus, if more breakthrough ideas means that the distribution changes in such a way that R(F) increases, examiners will be more lenient (since α^* will weakly increase) and applicants with ideas of value $v \leq \bar{v}$ will obfuscate more.

The next lemma shows that a first-order stochastic dominance implies an increase in $R(\cdot)$ and, therefore, leads to more obfuscation.

Proposition 4. *If F dominates G according to first-order stochastic dominance, then $R(F) \geq R(G)$. Hence, there is weakly more obfuscation under F than under G.*

We can also study how a change in the distribution of the value of ideas affects the examiner's equilibrium expected payoff, which is

$$U^*(F) = \begin{cases} (1 - F(\bar{v}))(1 - \pi(\sigma, 0))E[v|v > \bar{v}] & \text{, if } \alpha^*(F) < 1, \\ (1 - F(\bar{v}))E[v|v > \bar{v}] - c\pi(\sigma, e(1, \sigma))F(\bar{v}) & \text{, if } \alpha^*(F) = 1. \end{cases}$$

Let $F \succeq^{FOSD} G$. Directly from first-order stochastic dominance arguments we obtain $U^*(F) \geq U^*(G)$. Thus a "better" distribution makes the examiner weakly better off even though there is more obfuscation. Intuitively, this occurs because the loss from more obfuscation is smaller than the direct gain from an improvement in the distribution of values of ideas.

5 Investing in Knowledge

Good ideas are costly and applicants are more likely to find them if they invest in knowledge. Their incentives to invest depend on the payoff difference between having a good and a bad idea—determined by the subsequent level of obfuscation and the decision to grant patents by uncertain examiners—and on the *difficulty* of finding good ideas.

An applicant is characterized by her investment cost $k \sim F[0, A]$, with F continuously differentiable with density f. Each applicant observes k before investing to increase the probability of finding a good idea from μ_0 to μ_S. To avoid a large taxonomy of cases, in this section, we assume that $g_V = g_0 \equiv g$. In Appendix E, we analyze the case $g_V > g_0$, and show that additional parametric restrictions preserve the results in this section.

Definition 2. *An investment equilibrium is given by* (p^*, k^*) *such that*

1. *The proportion of good ideas in the economy is given by*

$$p^* = F(k^*)\mu_S + (1 - F(k^*))\mu_0 \tag{8}$$

2. *An applicant of type* k *invests if and only if* $k \leq k^*$*, where*

$$k^* = (\mu_S - \mu_0)[U_V(p^*) - U_0(p^*)] \tag{9}$$

is the type indifferent between investing and not investing, and $U_v(p)$ *is the expected payoff of an applicant with an idea of value* $v \in \{0, V\}$*, for any fixed* p*, with*

$$U_0(p) = \alpha^* \pi(\sigma, e^*)g - e^*, \tag{10}$$

$$U_V(p) = [1 - \pi(\sigma, 0) + \alpha^* \pi(\sigma, 0)]g. \tag{11}$$

Both the probability that an uncertain examiner grants a patent, α^*, and the level of obfuscation, e^*, depend on the proportion of good ideas, p^*. As shown in Section 3, e^* and α^* increase with p^*.

Proposition 5. *The difference* $U_V - U_0$ *is weakly decreasing in* p*.*

Figure 6 shows $U_V - U_0$ as a function of p, which is positive and non-increasing. With constrained examiners (Figure 6, left panel, $\sigma \leq \bar{\sigma}_1$), there is no obfuscation. Uncertain examiners either grant patents for sure or not at all, so $U_V - U_0$ is constant. With less constrained examiners (Figure 6, right panel, $\sigma > \bar{\sigma}_1$), a higher p increases α^*. In this case, applicants with bad ideas benefit more because they increase obfuscation, e^*, to create uncertainty for the examiner. Thus, the difference $U_V - U_0$ strictly decreases when $\alpha^* \in (0, 1)$.

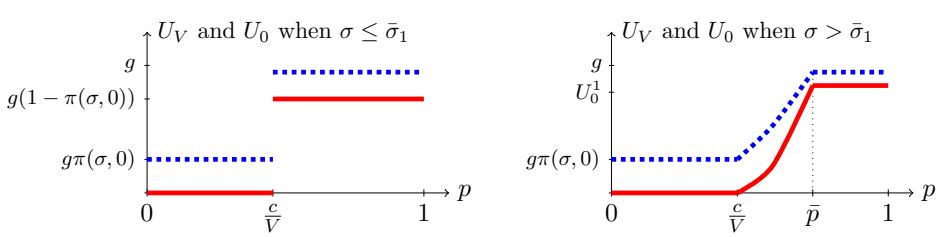

Figure 6: Applicant's equilibrium expected payoff as a function of p. The dashed line is U_V, and the solid line is U_0. Left panel: $\sigma \leq \bar{\sigma}_1$. Right panel: $\sigma > \bar{\sigma}_1$. We denote $U_0^1 = (1 - \pi(\sigma, e(1, \sigma)))g - e(1, \sigma)$.

Proposition 6. *There exists a unique equilibrium* (p^*, k^*) *that satisfies Definition 2.*

An equilibrium simultaneously determines the proportion of good ideas, p^*, and the investment, k^* (see Definition 2). Fixing everything else, the proportion of good ideas increases with investment (see Equation 8), whereas investment decreases with the proportion of good ideas (see Equation 9). Figure 7 illustrates the "supply" of good ideas for each investment level, $p(k)$, and the investment conditional on the proportion of good ideas, $k(p)$. The intersection of this two curves determines the equilibrium.

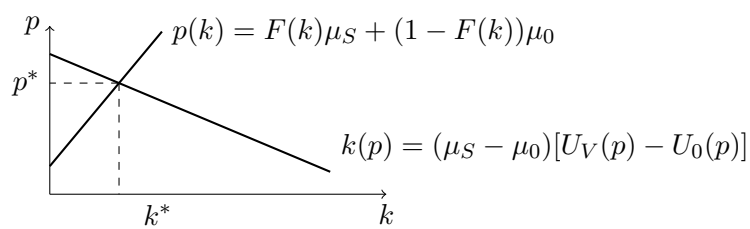

Figure 7: Proportion of good ideas and investment.

5.1 Resources for Examiners (Higher σ) and Private Returns (Higher g)

The direct effect of higher σ on investment is positive: more resourceful examiners find hard evidence more frequently so the difference $U_V - U_0$ increases. However, higher σ can increase or reduce obfuscation. Proposition 7 shows that the total effect of σ on investment is positive (i.e., the curve $k(p)$ in Figure 7 shifts to the right). There is no direct effect of σ on the supply of good ideas (i.e., the curve $p(k)$ in Figure 7 is independent of σ), so in equilibrium, both p^* and k^* increase with σ.

The private return, g, also impacts investment through the difference $U_V - U_0$ via two channels. First, it makes investing more attractive—good ideas are funded more often than bad ones—so g directly increases $U_V - U_0$. Second, higher g encourages applicants with bad ideas to obfuscate, directly increasing obfuscation, e^*, which reduces $U_V - U_0$. While in principle the effect of higher g on investment is ambiguous, we show that the aggregate effect is positive (i.e., higher g shifts the curve $k(p)$ in Figure 7 to the right).

Proposition 7. *p^* and k^* are weakly increasing in σ and g.*

The examiner, however, is not necessarily better off if σ or g goes up. In fact, Proposition 2 shows that the direct effect of a higher value of σ may be negative due to an increase in obfuscation effort. We must now add an extensive margin effect due to an increase in the

proportion of good ideas, which is always beneficial for the examiner:

$$\frac{du_S(\sigma, p^*(\sigma))}{d\sigma} = \underbrace{\frac{\partial u_S(\sigma, p^*(\sigma))}{\partial \sigma}}_{\text{Obfuscation}\leq 0} + \underbrace{\frac{\partial u_S(\sigma, p^*(\sigma))}{\partial p}}_{\text{More good ideas}\geq 0} \cdot \underbrace{\frac{\partial p^*(\sigma)}{\partial \sigma}}_{\text{Investment}\geq 0} \tag{12}$$

A higher σ makes the examiner worse off; the (negative) obfuscation term overcomes the positive effect of more good ideas. This condition is more demanding than (6). Thus, examiner's benefit from more resources is larger when the proportion of good ideas is endogenously determined. A similar argument shows that the examiner can be worse off when g goes up.

To shed light on allocative efficiency under different economic conditions, we study how exogenous changes to the economic environment affect the equilibrium proportion of good ideas, p^*, and the fraction of applicants that invest in knowledge, k^*.

5.2 Ideas are Getting Harder to Find

Recent literature has documented that ideas are getting harder to find (Jones, 2009; Bloom et al., 2020), which has a direct impact on applicants' investment in knowledge and also on their ex-post behavior (obfuscation).

There are three channels through which ideas might be harder to find: (1) a reduction in "low-hanging fruit" (a decrease in μ_0); (2) a reduction in the returns to knowledge (a decrease in μ_S); or (3) an increase in the the burden of knowledge (a FOSD increase in F). A reduction in low-hanging fruit can be interpreted as a field reaching "maturity," so applicants are less likely to come across good ideas without investing in knowledge. A reduction in returns to knowledge can be interpreted as a more difficult field even for specialists.[10] A FOSD increase of F can be interpreted as costlier acquisition of knowledge—reaching the frontier of knowledge takes longer and scientists' breakthroughs occur later in life. Each of these channels impact the equilibrium differently.

Proposition 8. *We have:*

1. *If μ_0 decreases, k^* increases; p^* (and obfuscation) decreases if and only if*

$$\frac{1 - F(k^*)}{f(k^*)} > k^*. \tag{13}$$

[10] Arguably, this is the current situation in the field of theoretical physics (see, e.g., Butterworth, 2019).

2. *If μ_S decreases, p^* (and obfuscation) decreases; k^* decreases if and only if*

$$k^* \geq -(\mu_S - \mu_0)^2 \left(\frac{\partial [U_V(p^*) - U_0(p^*)]}{\partial p} \right) F(k^*). \tag{14}$$

3. *If F increases in the FOSD order, then p^* (and obfuscation) and k^* decrease.*

First, a reduction of low-hanging fruit, μ_0, leads to higher aggregate investment (i.e., $k(p)$ shifts to the right in Figure 7). An applicant who does not invest has a lower chance of receiving U_V, motivating investment. However, lower μ_0 also directly decreases the supply of good ideas (i.e., $p(k)$ shifts down in Figure 7). Thus, k^* always increase and while more applicants invest in knowledge, those who do not invest are less likely to find good ideas, so the impact on the proportion of good ideas, p^* is ambiguous. If p^* decreases, obfuscation decreases and $U_V - U_0$ increases.

Second, a reduction in the productivity of knowledge, μ_S, has an ambiguous effect on k^*. The direct effect of lowering μ_S is to decrease investment because its return is lower (i.e., $k(p)$ shifts to the left in Figure 7) and to increase the supply of good ideas (i.e., $p(k)$ shifts up in Figure 7). Fewer applicants invest, and those who do are less likely to find good ideas, therefore, p^* decreases.

Third, a FOSD increase in F means that investing is costlier for more applicants. This only impacts (negatively) the the supply of good ideas (i.e., $p(k)$ shifts down in Figure 7), so k^* and p^* decrease.

6 Discussion

Many economic situations involve a privately-informed sender who seeks approval from a receiver whose objective is to approve only "high types": Inventors want all their inventions patented, but the patent office wants to grant patents for novel ideas only; Firms want investor to invest in them, but investors want to invest only in financially-sound firms; Academics want their papers published, but journals want to publish good papers only; Scientists want to receive grant funding, but sponsors want to fund good ideas only. Obfuscation—the effort to disrupt comprehension—has been empirically documented in all these instances. Obfuscation is problematic since it reduces allocative efficiency ex-post and dampens investment incentives ex-ante.

We study strategic obfuscation in patent applications. By reading a patent application, the examiner may find concrete evidence about an idea's value. Without concrete evidence, the examiner must decide to grant a patent or reject the application. Obfuscation (or lack of clarity) arises when examiners are lenient. We show that granting patents with higher probability induces more obfuscation (Proposition 1 and Figure 6), and in Section 4.1, we show that leniency induces less clarity. In subsection 3.1, we use data from US patent applications to empirically assess the validity of this conclusion.

A possible solution is to provide more resources to examiners. We show that such a policy could backfire. With more resourceful examiners, patent applicants with bad ideas obfuscate more, since their only chance to get a patent is to leave the examiner uncertain about the idea's value after the examination. As a consequence, we show that the examiner's payoff is non-monotonic in the amount of resources (Proposition 2). We show that this result is robust to a model where applicants can obfuscate and clarify, and where there is a continuum of types of ideas (Section 4).

We also study the impact of obfuscation on ex-ante investment in knowledge, which determines the equilibrium proportion of good ideas in the economy. This determines the amount of obfuscation effort and, hence, the equilibrium payoffs of applicants with good and bad ideas (Proposition 6). We show that the impact of "science getting harder" on obfuscation is ambiguous: a reduction of low-hanging fruit may increase or decrease obfuscation, but obfuscation always decreases when knowledge is less productive or costlier to acquire (Proposition 8). Furthermore, even after endogenizing the proportion of good ideas, more resources for examiners can backfire Proposition 7).

References

Ashtor, Jonathan H (2022) "Modeling patent clarity," *Research Policy*, Vol. 51, p. 104415.

Atal, Vidya and Talia Bar (2010) "Prior art: To search or not to search," *International Journal of Industrial Organization*, Vol. 28, pp. 507–521.

Azoulay, Pierre and Danielle Li (2020) "Scientific Grant Funding,"Technical report, National Bureau of Economic Research.

Bar, Talia and Sidartha Gordon (2014) "Optimal project selection mechanisms," *American Economic Journal: Microeconomics*, Vol. 6, pp. 227–55.

Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb (2020) "Are ideas getting harder to find?" *American Economic Review*, Vol. 110, pp. 1104–44.

Boleslavsky, Raphael, Bruce Ian Carlin, and Christopher Cotton (2019) "Disincentive effects of evaluation,"Technical report, Queen's Economics Department Working Paper.

Boleslavsky, Raphael and Christopher Cotton (2018) "Limited capacity in project selection: Competition through evidence production," *Economic Theory*, Vol. 65, pp. 385–421.

Bushee, Brian J, Ian D Gow, and Daniel J Taylor (2018) "Linguistic complexity in firm disclosures: Obfuscation or information?" *Journal of Accounting Research*, Vol. 56, pp. 85–121.

Butterworth, Jon (2019) "A struggle for the soul of theoretical physics," *Nature*, Vol. 568, pp. 309–309.

Cao, Sean, Wei Jiang, Baozhong Yang, Alan L Zhang et al. (2020) "How to Talk When a Machine is Listening: Corporate Disclosure in the Age of AI,"Technical report, National Bureau of Economic Research.

Che, Yeon-Koo, Wouter Dessein, and Navin Kartik (2013) "Pandering to persuade," *American Economic Review*, Vol. 103, pp. 47–79.

Chubb, Jennifer and Richard Watermeyer (2017) "Artifice or integrity in the marketization of research impact? Investigating the moral economy of (pathways to) impact statements within research funding proposals in the UK and Australia," *Studies in Higher Education*, Vol. 42, pp. 2360–2372.

de Clippel, Geoffroy and Kareen Rozen (2020) "Communication, Perception, and Strategic Obfuscation,"Technical report, Working Paper.

Crawford, Vincent P and Joel Sobel (1982) "Strategic information transmission," *Econometrica: Journal of the Econometric Society*, pp. 1431–1451.

Dewatripont, Mathias and Jean Tirole (2005) "Modes of communication," *Journal of political economy*, Vol. 113, pp. 1217–1238.

Drake, Michael S, Jeffrey Hales, and Lynn Rees (2019) "Disclosure Overload? A Professional User Perspective on the Usefulness of General Purpose Financial Statements," *Contemporary Accounting Research*, Vol. 36, pp. 1935–1965.

Ellison, Glenn and Alexander Wolitzky (2012) "A search cost model of obfuscation," *The RAND Journal of Economics*, Vol. 43, pp. 417–441.

Frakes, Michael D and Melissa F Wasserman (2017) "Is the time allocated to review patent applications inducing examiners to grant invalid patents? Evidence from microlevel application data," *Review of Economics and Statistics*, Vol. 99, pp. 550–563.

Frankel, Alex and Navin Kartik (2019) "Muddled information," *Journal of Political Economy*, Vol. 127, pp. 1739–1776.

Freilich, Janet (2017) "Patent Clutter," *Iowa L. Rev.*, Vol. 103, p. 925.

Gu, Yiquan and Tobias Wenzel (2014) "Strategic obfuscation and consumer protection policy," *The Journal of Industrial Economics*, Vol. 62, pp. 632–660.

Head, Megan L, Luke Holman, Rob Lanfear, Andrew T Kahn, and Michael D Jennions (2015) "The extent and consequences of p-hacking in science," *PLoS Biol*, Vol. 13, p. e1002106.

Howell, Sabrina T (2017) "Financing innovation: Evidence from R&D grants," *American Economic Review*, Vol. 107, pp. 1136–64.

Iliev, Rumen and Robert Axelrod (2017) "The paradox of abstraction: Precision versus concreteness," *Journal of psycholinguistic research*, Vol. 46, pp. 715–729.

Jin, Ginger Zhe, Michael Luca, and Daniel Martin (2021) "Complex disclosure," *Management Science*.

Jones, Benjamin F (2009) "The burden of knowledge and the "death of the renaissance man": Is innovation getting harder?" *The Review of Economic Studies*, Vol. 76, pp. 283–317.

Kartik, Navin (2009) "Strategic communication with lying costs," *The Review of Economic Studies*, Vol. 76, pp. 1359–1395.

Kartik, Navin, Marco Ottaviani, and Francesco Squintani (2007) "Credulity, lies, and costly talk," *Journal of Economic theory*, Vol. 134, pp. 93–116.

Kuhn, Jeffrey, Kenneth Younge, and Alan Marco (2020) "Patent citations reexamined," *The RAND Journal of Economics*, Vol. 51, pp. 109–132.

Langinier, Corinne and Philippe Marcoul (2020) "Monetary and implicit incentives of patent examiners," *Journal of Economics and Business*, p. 105906.

Larrimore, Laura, Li Jiang, Jeff Larrimore, David Markowitz, and Scott Gorski (2011) "Peer to peer lending: The relationship between language features, trustworthiness, and persuasion success," *Journal of Applied Communication Research*, Vol. 39, pp. 19–37.

Legros, Patrick and Andrew F Newman (2002) "Courts, contracts, and interference," *European Economic Review*, Vol. 46, pp. 734–744.

Li, Danielle (2017) "Expertise versus Bias in Evaluation: Evidence from the NIH," *American Economic Journal: Applied Economics*, Vol. 9, pp. 60–92.

Li, Feng (2008) "Annual report readability, current earnings, and earnings persistence," *Journal of Accounting and economics*, Vol. 45, pp. 221–247.

Markowitz, David M (2019) "What words are worth: National Science Foundation grant abstracts indicate award funding," *Journal of Language and Social Psychology*, Vol. 38, pp. 264–282.

Markowitz, David M and Jeffrey T Hancock (2016) "Linguistic obfuscation in fraudulent science," *Journal of Language and Social Psychology*, Vol. 35, pp. 435–445.

Milgrom, Paul and John Roberts (1986) "Relying on the information of interested parties," *The RAND Journal of Economics*, pp. 18–32.

Nosek, Brian A, Jeffrey R Spies, and Matt Motyl (2012) "Scientific utopia: II. Restructuring incentives and practices to promote truth over publishability," *Perspectives on Psychological Science*, Vol. 7, pp. 615–631.

Perez-Richet, Eduardo and Delphine Prady (2012) "Complicating to persuade?".

Persson, Petra (2018) "Attention manipulation and information overload," *Behavioural Public Policy*, Vol. 2, pp. 78–106.

Stahl, Konrad and Roland Strausz (2017) "Certification and market transparency," *The Review of Economic Studies*, Vol. 84, pp. 1842–1868.

Vinkers, Christiaan H, Joeri K Tijdink, and Willem M Otte (2015) "Use of positive and negative words in scientific PubMed abstracts between 1974 and 2014: retrospective analysis," *Bmj*, Vol. 351, p. h6467.

A Appendix: Proofs

Proof of Proposition 1

Proof. 1. and 2. are straightforward. For 3. the sponsor is indifferent, so any $\alpha \in [0,1]$ is optimal. However, α^* must be consistent with Q^* according to Bayes rule.

We now characterize the equilibrium as a function of p and σ.

a. Equilibria without obfuscation. Let $\bar{\sigma}_1(g)$ be such that $e(1,\sigma,g) = 0$, for all $\sigma \leq \bar{\sigma}_1(g)$.

1. When $p(V+c) < c$, then $Q^*(V+c) < c$, because $Q^* \leq p$, so $\alpha^* = 0$ and $e^* = 0$ is the unique PBE.

2. When $p(V+c) = c$, and $\sigma > \bar{\sigma}_1$, we must have $\alpha^* = 0$. If, by contradiction, $\alpha^* > 0$, then $e^* > 0$ and $Q^* < p$, so $Q^*(V+c) < c$ which would require $\alpha^* = 0$. If $p(V+c) = c$, and $\sigma \leq \bar{\sigma}_1$, then $e^* = 0$ and $Q^* = p$, regardless of the value of $\alpha^* \in [0,1]$.

3. When $p(V+c) > c$ and $\sigma \leq \bar{\sigma}_1$, Note that for any $\alpha^* \in [0,1]$, $0 \leq e^* = e(\alpha^*,\sigma) \leq e(1,\sigma) = 0$ because $e(\cdot,\sigma)$ is weakly increasing. Thus, in any PBE we must have $Q^* = p$ and $\alpha^* = 1$.

b. Equilibria with obfuscation. These equilibria occurs when $p(V+c) > c$ and $\sigma > \bar{\sigma}_1$. We begin with a preliminary result.

Claim: There exists a threshold $\bar{\sigma}_2 \in (\bar{\sigma}_1, 1)$ such that inequality (3) holds when $e^* = e(1,\sigma)$ if and only if $\sigma \leq \bar{\sigma}_2$. Furthermore, $\bar{\sigma}_2$ increases in p and decreases in g. We denote by $\bar{p}(\sigma)$ the inverse of $\bar{\sigma}_2(p)$.

Proof of the Claim: Let $f(\sigma) = \pi(\sigma, 0)pV - \pi(\sigma, e^*)(1-p)c$, where $e^* = e(1, \sigma)$. For $\sigma \leq \bar{\sigma}_1$, $e(1, \sigma, g) = 0$ and, therefore, $f(\sigma) = \pi(\sigma, 0)(p(V + c) - c) > 0$. For $\sigma > \bar{\sigma}_1$, $e(1, \sigma) > 0$, and

$$f'(\sigma) = \pi_\sigma(\sigma, 0)pV - \pi_\sigma(\sigma, e^*)(1-p)c - \pi_e(\sigma, e^*)e_\sigma(1, \sigma)(1-p)c.$$

π_σ is increasing in e, by the supermodularity of $\pi(\cdot)$, so $\pi_\sigma(\sigma, e^*) \geq \pi_\sigma(\sigma, 0)$, which implies

$$f'(\sigma) < \pi_\sigma(\sigma, 0)(p(V + c) - c) - \pi_e(\sigma, e(1, \sigma))e_\sigma(1, \sigma, g)(1-p)c < 0$$

The last inequality uses that $\pi_\sigma(\sigma, 0)(p(V + c) - c) < 0$, $\pi_e > 0$, and $e_\sigma \geq 0$. Finally, using that $\pi(0, e) = 1$, we get $f(0) = p(V + c) - c > 0$, and using that $0 = \pi(1, 0) < \pi(1, e^*)$ we get $f(1) < 0$. Thus, there exists a unique threshold $\bar{\sigma}_2$ such that $f(\bar{\sigma}_2) = 0$.

The comparative statics are trivial. Consider any parameter θ. We have $f(\theta, \bar{\sigma}_2) = 0$. Then,

$$\frac{\partial f}{\partial \theta} + \frac{\partial f}{\partial \sigma}\frac{\partial \bar{\sigma}_2}{\partial \theta} = 0.$$

Since f is strictly decreasing in σ, we have sign $\frac{\partial \bar{\sigma}_2}{\partial \theta}$ = sign $\frac{\partial f}{\partial \theta}$. [*End of the Claim*]

We have two cases to consider:

1. When $p(V + c) > c$ and $\sigma \in (\bar{\sigma}_1, \bar{\sigma}_2)$, we must have $Q^*(V + c) > c$. Otherwise, if $Q^*(V + c) < c$ we have $\alpha^* = 0$, $e^* = 0$, which would imply that $Q^* = p$. But that would be a contradiction, because $c < p(V + c) = Q^*(V + c) < c$. Next, if $Q^*(V + c) = c$, then we need to find $\alpha^* \in [0, 1]$ consistent with Equation 17. But this is impossible, because $f(\cdot)$ (defined in the proof of the Claim) is such that $f(\sigma) > 0$ for $\sigma < \bar{\sigma}_2$, so no $\alpha \in [0, 1]$ solves Equation 17. Therefore, it must be that if $Q^*(V + c) > c$, so $\alpha^* = 1$ and $e^* = e(1, \sigma, g)$, which is consistent when $\sigma < \bar{\sigma}_2$ and Q^* satisfies Equation 2.

2. When $p(V + c) > c$ and $\sigma \geq \bar{\sigma}_2$, there is no equilibrium with $Q^*(V + c) < c$. Otherwise, we would have $\alpha^* = 0$, $e^* = 0$ and $Q^* = p$, which is a contradiction. Also we cannot have $Q^*(V + c) > c$ because we would have $\alpha^* = 1$ and $e^* = e(1, \sigma)$, which imply $Q^*(V + c) \leq c$ (by the definition of $\bar{\sigma}_2$). Thus, we must have $Q^*(V + c) = c$. For any $\sigma \geq \bar{\sigma}_2$, when $\alpha^* = 0$, we get $Q^* = p(V + c) > c$, and when $\alpha^* = 1$ we get $Q^*(V + c) < c$. Then, by continuity, we can always find $\alpha(\sigma, p, g) \in [0, 1]$ such that Equation 17 holds. By continuity, $\alpha(\bar{\sigma}_2, p, g) = 1$ and $\alpha(1, p, g) = 0$. The monotonicity in α is direct from Equation 17.

□

Proof of Proposition 2

Proof. The examiner's equilibrium expected payoff is

$$u_E = (1 - \pi(\sigma, 0))pV + \alpha^*[\pi(\sigma, 0)pV - \pi(\sigma, e^*)(1 - p)c].$$

The sponsor mixes only when she is indifferent, so $\alpha^* \in (0, 1)$ only when $Q^*(V + c) = c$ or, equivalently, $\pi(\sigma, 0)pV = \pi(\sigma, e^*)(1-p)c$. In that case, $u_E = (1-\pi(\sigma, 0))pV$, which is strictly increasing in σ. When $\alpha^* = 1$ we have $u_E = pV - \pi(\sigma, e(1, \sigma))(1-p)c$. Thus, u_E decreases in σ only when $\pi(\sigma, e(1, \sigma))$ increases, which is the inequality in the proposition. □

Proof of Proposition 3

Proof. To see why, let

$$K(\alpha) = \frac{\pi(\sigma, e(\alpha, \sigma))}{\pi(\sigma, 0)}c.$$

Note that $K(\cdot)$ is non-decreasing, since α weakly increases obfuscation effort, which weakly increases the probability of leaving the examiner uncertain. Using $R(F)$ and $K(\alpha)$ we can re-characterize the equilibrium: $\alpha^* = 1$ if $R(F) \geq K(1)$; $\alpha^* = 0$ if $R(F) \leq K(0)$; $\alpha^* \in [0, 1]$, where α^* is pinned down by $R(F) = K(\alpha^*)$.

For two distributions F and G, whenever $R(F) \geq R(G)$, uncertain examiners are more likely to grant a patent when ideas comes from F than when they come from G. As a consequence, obfuscation is weakly higher for a distribution F than G when $R(F) \geq R(G)$. □

Proof of Proposition 4

Proof. We just need to show that $\mathrm{R}(F) \geq \mathrm{R}(G)$ whenever $F \succeq^{FOSD} G$. That is,

$$\frac{1}{F(\bar{v})}\int_v^\infty v dF(v) \geq \frac{1}{G(\bar{v})}\int_v^\infty v dG(v).$$

We know that $F(\bar{v}) \leq G(\bar{v})$ whenever $F \succeq^{FOSD} G$. Also, since $u(v) = v1(v \geq \bar{v})$ is a non-decreasing function, first-order stochastic dominance implies have

$$\int_{-\infty}^\infty u(v)dF(v) \geq \int_{-\infty}^\infty u(v)dG(v).$$

Combining these two inequalities we get $R(F) \geq R(G)$. □

Proof of Proposition 5

Proof. We have $U_V(p) - U_0(p) = g(1-\pi(\sigma,0)) + g\alpha^*\pi(\sigma,0) - g\alpha^*\pi(\sigma,e^*) + e^*$. This difference is always positive because from optimality

$$g\alpha^*\pi(\sigma,0) \leq g\alpha^*\pi(\sigma,e^*) - e^*.$$

Also note that for any p such that $e^* = 0$, $U_V(p) - U_0(p) = g(1-\pi(\sigma,0))$, and for any p such that $\alpha^* = 1$, $U_V(p) - U_0(p) = g - g\pi(\sigma, e(1,\sigma)) - e(1,\sigma)$. More generally, $U_V - U_0$ can be written as

$$(U_V - U_0)(p) = g\pi(\sigma,0) + V(p,0) - V(p, e(p,\sigma,g)),$$

where $V(p,e) = g\alpha(p,\sigma,g)(1-\pi(\sigma,e)) - e$. Taking the derivative with respect to p and using the envelope theorem we get

$$(U_V - U_0)'(p) = g\frac{\partial\alpha(p,\sigma,g)}{\partial p}[\pi(\sigma,e^*) - \pi(\sigma,0)] \leq 0.$$

□

Proof of Proposition 6

Proof. To simplify the notation, we call α the solution to (17), e the solution to (1). An equilibrium (p,k) solves:

$$\begin{aligned} p &= \mu_0 + F(k)(\mu_S - \mu_0) \\ k &= (\mu_S - \mu_0)[U_V(p) - U_0(p)] \end{aligned}$$

We define $H(p) = \mu_0 + F((\mu_S - \mu_0)[U_V(\sigma,p,g) - U_0(\sigma,p,g)])(\mu_S - \mu_0)$. In equilibrium, p solves the fixed-point equation $p = H(p)$.

1. Existence: $H(\cdot)$ is continuous and $H(p) \in [0,1]$. By the intermediate value theorem, an equilibrium exists.

2. Uniqueness: We have $H(p) = H(0)$ for any $p \leq c/(V+c)$, because $\alpha^* = 0$ and $e^* = 0$. Also, $H(p) = H(1)$ for any $p > \bar{p}(\sigma)$, because $\alpha^* = 1$ and $e^* = e(1,\sigma)$. By optimality $g(1-\pi(\sigma,e(1,\sigma))) - e(1,\sigma) \geq g(1-\pi(\sigma,0))$, so $H(0) \geq H(1)$. Finally, $H(\cdot)$ is decreasing in p, for any $p \in (c/(V+c), \bar{p})$ because $F(\cdot)$ is increasing and $U_V - U_0$ is decreasing in p (see Proposition 5). Thus, the solution to $p = H(p)$ is unique. □

Proof of Proposition 8

Proof. Consider the equilibrium conditions:

$$\begin{aligned} p^* &= \mu_0 + F(k^*)(\mu_S - \mu_0) \\ k^* &= (\mu_S - \mu_0)[U_V(\sigma, p^*, g) - U_0(\sigma, p^*, g)] \end{aligned}$$

1. Differentiating the equilibrium conditions with respect to μ_0, denoted by $()'$:

$$\begin{aligned} (p^*)' &= 1 - F(k^*) + (\mu_S - \mu_0) f(k^*)(k^*)' \\ (k^*)' &= \frac{-k^*}{\mu_S - \mu_0} + (\mu_S - \mu_0) \frac{\partial[U_V(\sigma, p^*, g) - U_0(\sigma, p^*, g)]}{\partial p}(p^*)' \end{aligned}$$

Replacing $(k^*)'$ from the second equation in the first equation we get

$$(p^*)' = 1 - F(k^*) - f(k^*)k^* + (\mu_S - \mu_0)^2 \frac{\partial[U_V(\sigma, p^*, g) - U_0(\sigma, p^*, g)]}{\partial p}(p^*)'$$

Given that $\frac{\partial[U_V(\sigma,p^*,g)-U_0(\sigma,p^*,g)]}{\partial p} \leq 0$, the sign of $(p^*)'$ is the same sign of $1 - F(k^*) - f(k^*)k^*$.

If $(p^*)' > 0$, it is immediate from the second equilibrium condition that $(k^*)' < 0$. If $(p^*)' < 0$, then from the first equation it must be that $(k^*)' < 0$.

2. Differentiating the equilibrium conditions with respect to μ_S, denoted by $()'$:

$$\begin{aligned} (p^*)' &= F(k^*) + (\mu_S - \mu_0) f(k^*)(k^*)' \\ (k^*)' &= \frac{k^*}{\mu_S - \mu_0} + (\mu_S - \mu_0) \frac{\partial[U_V(\sigma, p^*, g) - U_0(\sigma, p^*, g)]}{\partial p}(p^*)' \end{aligned}$$

Replacing the second equation in the first and using that $\frac{\partial[U_V(\sigma,p^*,g)-U_0(\sigma,p^*,g)]}{\partial p} \leq 0$, we get that $(p^*)' > 0$. Replacing $(p^*)'$ from the first equation into the second equation we get that

$$\text{sign } (k^*)' = \text{ sign } \left[k^* + (\mu_S - \mu_0)^2 \left(\frac{\partial[U_V(\sigma, p^*, g) - U_0(\sigma, p^*, g)]}{\partial p} \right) F(k^*) \right]$$

3. Suppose by that $p^*(F) > p^*(G)$. Then, using that $U_V - U_0$ is decreasing in p we would have $k^*(F) \leq k^*(G)$. Using the definition of p^*, that $F(\cdot)$ is weakly increasing in that $F(s) \leq G(s)$ for all s we have the following contradiction:

$$p^*(F) = \mu_0 + F(k^*(F))(\mu_S - \mu_0) \leq \mu_0 + F(k^*(G))(\mu_S - \mu_0) \leq \mu_0 + G(k^*(G))(\mu_S - \mu_0) = p^*(G).$$

□

Proof of Proposition 7

Proof. In equilibrium we have $p^* = H(g, p^*)$, so

$$\frac{\partial p}{\partial g} = \frac{\partial H}{\partial g} + \frac{\partial H}{\partial p}\frac{\partial p}{\partial g} \iff \frac{\partial p}{\partial g} = \frac{\frac{\partial H}{\partial g}}{1 - \frac{\partial H}{\partial p}}.$$

Given that $\frac{\partial H}{\partial p} \leq 0$ (see the proof of Proposition 6) we have:

$$\text{sign}\,\frac{\partial p}{\partial g} = \text{sign}\,\frac{\partial H}{\partial g} = \text{sign}\,\frac{\partial k}{\partial g} = \frac{\partial (U_V - U_0)}{\partial g}.$$

Taking the derivative with respect to g and using the envelope theorem we get

$$(U_V - U_0)'(g) = \alpha^* \pi(\sigma, e^*) + (1 - \alpha^*)\pi(\sigma, 0) + g\frac{\partial \alpha^*}{\partial g}[\pi(\sigma, e^*) - \pi(\sigma, 0)] \geq 0,$$

because $\frac{\partial \alpha(p,\sigma,g)}{\partial g} \leq 0$ and $\pi(\sigma, e^*) \leq \pi(\sigma, 0)$, so k^* and p^* are (weakly) increasing in g.

Analogous to the previous case we have:

$$\text{sign}\,\frac{\partial p}{\partial \sigma} = \text{sign}\,\frac{\partial H}{\partial \sigma} = \text{sign}\,\frac{\partial k}{\partial \sigma} = \frac{\partial (U_V - U_0)}{\partial \sigma}.$$

We also have $(U_V - U_0)(\sigma) = g\pi(\sigma, 0) + V(\sigma, 0) - V(\sigma, e(p, \sigma, g))$, where $V(\sigma, e) = g\alpha(p, \sigma, g)(1 - \pi(\sigma, e)) - e$. Taking derivative with respect to σ and using the envelope theorem we get

$$(U_V - U_0)'(\sigma) = g\left[\alpha\pi_\sigma(\sigma, e^*) + (1 - \alpha)\pi_\sigma(\sigma, 0) + \frac{\partial \alpha(p, \sigma, g)}{\partial \sigma}[\pi(\sigma, e^*) - \pi(\sigma, 0)]\right] \geq 0,$$

because $\frac{\partial \alpha(p,\sigma,g)}{\partial \sigma} \leq 0$ and $\pi(\sigma, e^*) \leq \pi(\sigma, 0)$. Therefore, k^* and p^* are non-decreasing in σ. $\square$

Online Appendix

The Benefit of the Doubt: Patent Examination Under Strategic Obfuscation

Nicolás Figueroa and Jorge Lemus

Supplemental Material – Intended for Online Publication

A Elastic Supply of Examiner Effort

In this section we study the examiner's strategic provision of effort. The examiner's payoff in the baseline model corresponds to the patent office's payoff. In this section, we separate the patent office and the examiner, and we assume the patent office provides an incentive contract to the examiner who chooses effort based on this contract.

In order to provide incentives, the patent office pays the examiner an amount t for finding hard evidence. Effort is costly and the cost is given by $\tilde{c}(\sigma)$ with $\tilde{c}_\sigma > 0$ and $\tilde{c}_{\sigma\sigma} > 0$. The examiner plays a simultaneous game with the applicant. For a given obfuscation effort e, the reviewer chooses effort σ to solve

$$\max_{\sigma \in [0,1]} t[p(1-\pi(\sigma,0)) + (1-p)(1-\pi(\sigma,e))] - \tilde{c}(\sigma).$$

Defining $c(\sigma,t) = \frac{\tilde{c}(\sigma)}{t}$ we can write the problem as

$$\max_{\sigma \in [0,1]} [p(1-\pi(\sigma,0)) + (1-p)(1-\pi(\sigma,e))] - c(\sigma,t). \tag{15}$$

The baseline case corresponds to $c(\sigma,t) = 0$ for $\sigma \leq t$ and $c(\sigma,t) = +\infty$ for $\sigma > t$, and therefore the reviewer chooses $\sigma = t$ against any effort e, since it is a dominant strategy, and we identified $\sigma = t \in [0,1]$.

By the supermodularity assumption, $\pi_{\sigma,e} > 0$, the reviewer's best response, $\sigma_{BR}(e,p,t)$ (the solution to 15), is decreasing in e, increasing in p, and increasing in t (since $c_{\sigma,t}(\sigma,t) < 0$). The applicant chooses obfuscation as in the baseline case: $e(\alpha,\sigma)$.

It is easy to verify that $\sigma_{BR}(e(\alpha,\sigma^*),p,t)$ is uniquely determined. The equilibrium definition is only slightly modified by the endogeneity of σ.

Definition 3. *An equilibrium is characterized by* $(Q^*, \alpha^*, e^*, \sigma^*)$ *that satisfy:*

1. *The applicant's obfuscation effort is chosen optimally:* $e^* = e(\alpha^*, \sigma^*, g)$
2. *The reviewer's examination effort is chosen optimally:* $\sigma^* = \sigma_{BR}(e^*, p, t)$
3. *Beliefs* Q^* *are consistent Bayes' rule:*

$$Q^* = \frac{\pi(\sigma^*,0)p}{\pi(\sigma^*,0)p + \pi(\sigma^*,e^*)(1-p)}, \tag{16}$$

4. *The funding decision is consistent with the sponsor's posterior belief:*

$$\alpha^* \in \arg\max_{\alpha \in [0,1]} \alpha \cdot (Q^* V + (1 - Q^*)(-c)).$$

As in the baseline model, an uncertain examiner's benefit of granting a patent is $B = (1 - \pi(\sigma^*, 0))pV$ and its cost is $C = (1 - \pi(\sigma^*, e^*))(1 - p)c$. Now, both σ^* and e^* depend on p. Therefore the benefit is not always was strictly increasing in p, since σ^* increases, which can decrease B. The benefit is less than the cost ($B < C$) iff $Q^*(V + c) < c$, which allows us to write an analoguos version of Proposition 1.

Proposition 9. *Let Q^* be the examiner's posterior belief. In equilibrium we have:*

1. *$Q^*(V + c) > c$, an uncertain examiner grants a patent ($\alpha^* = 1$); applicants obfuscation and reviewer's effort are given by $e^* = e(1, \sigma^*)$ and $\sigma^* = \sigma_{BR}(e^*, p, t)$.*

2. *$Q^*(V + c) < c$, an uncertain examiner rejects the application ($\alpha^* = 0$); applicants with bad ideas do not obfuscate, $e^* = 0$, and reviewer's effort is given by $\sigma^* = \sigma_{BR}(0, p, t)$.*

3. *If $Q^*(V + c) = c$, an uncertain examiner is indifferent between granting a patent or rejecting the application. The examiner grants a patent with probability $\alpha^* \in [0, 1]$ and applicants with bad ideas obfuscate by $e^* = e(\alpha^*, \sigma^*)$ and reviewers choose $\sigma^* = \sigma_{BR}(e^*, p, t)$. The probability α^* is pinned down by Bayes rule and satisfies*

$$\pi(\sigma^*, 0)pV = \pi(\sigma^*, e(\alpha^*, \sigma^*))(1 - p)c. \tag{17}$$

Also analogously to the baseline model, it is possible to prove and characterize existence for any set of parameters. We omit the proof, which it is quite similar to the proof of Proposition 1. Instead, by means of an example, we show the qualitative features of the equilibrium.

Consider $\pi(\sigma, e) = 1 - \sigma q(e)$, where $q(e) = \max\{0, 1 - M\sqrt{e}\}$, and $c(\sigma, t) = \frac{\sigma^2}{2t}$. Under these assumptions, for any $\sigma > 0$ the applicant always exerts some obfuscation effort because $q'(e) \to -\infty$ as $e \to 0$. It is easy to show that the best responses of the applicant and the examiner are, respectively,

$$e_{BR}(\sigma) = \min\left\{\left(\frac{M\alpha g\sigma}{2}\right)^2, \frac{1}{M^2}\right\}, \quad \sigma_{BR}(e, p, t) = \min\{(p + (1 - p)q(e))t, 1\}. \tag{18}$$

As an illustration, consider the parameters $V = 5$, $c = 1$, $M = 1.5$, $t = 2/3$, $\alpha = g = 1$. To study the PBE equilibrium, we plot the benefit and the cost for two cases: $\alpha^* = 1$ and

$\alpha^* = 0$. If we set $\alpha^* = 1$, Figure A.1 (left panel) shows that the benefit outweighs the cost for $p \geq 0.25$. Thus, it is an equilibrium to set $\alpha^* = 1$ for any $p > 0.25$. If we set $\alpha^* = 0$, Figure A.1 (right panel) shows that the benefit outweighs the cost for $p \leq 0.2$. Thus, it is an equilibrium to set $\alpha^* = 0$ for any $p < 0.2$. For the intermediate range, $p \in (0.2, 0.25)$, there is no pure-strategy equilibrium, and the examiner chooses α^* to equate the benefit and the cost. The threshold 0.2 corresponds to c/V and the threshold 0.25, which corresponds to $\bar{p}(\gamma)$ in the main text, is determined by the parameters t, V, c, and g.

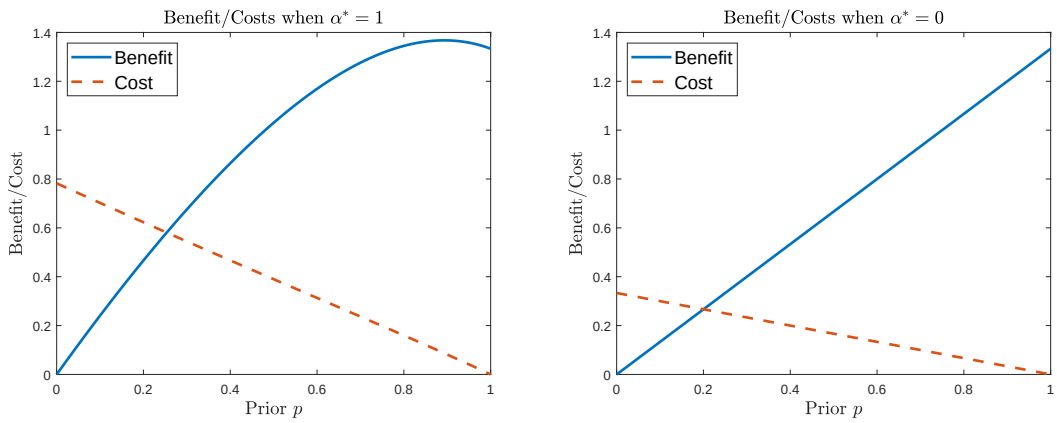

Figure A.1: Benefit (solid line) and cost (dashed line) for $\alpha^* = 1$ (left) and $\alpha^* = 0$ (right).

Figure A.2 shows the taxonomy of equilibria for different values of p and different types of examiners. The left panel shows the equilibria in baseline case, where the cost function is $c(\sigma, t) = 0$ for $\sigma \leq t$ and $+\infty$ for $\sigma > t$ (this corresponds to Figure 2; with this obfuscation technology, $\bar{\sigma}_1 = 0$). The right panel shows the equilibria when the cost function is $c(\sigma, t) = \frac{\sigma^2}{2t}$. Qualitatively, the equilibrium taxonomy is preserved.

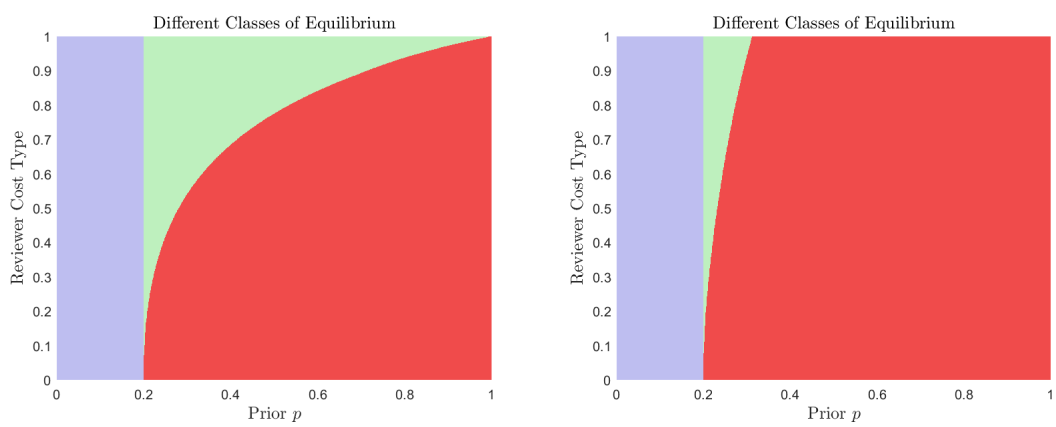

Figure A.2: Different regions of equilibrium

B Intermediate Signals

In the baseline model, the examiner either learns the value of the idea or is uncertain after examination. In this section, we augment our baseline framework to accommodate for different degrees of uncertainty after examination. Here, the examiner either learns the value of the idea or receives a noisy signal about the quality of the proposal. That is, the examiner either observes the true value of the idea, $v \in \{0, V\}$, or a *signal* $s \in S$ that conveys "soft evidence" (inconclusive). In contrast to the baseline model, the signal is positively *correlated* with the value of the idea, so good ideas are more likely to produce higher signals. Specifically, $s \sim F(s|v)$ where $f(s|v)$ is the density of F and satisfies the monotone-likelihood ratio (MLR) property holds, i.e., $\frac{f(s|V)}{f(s|0)}$ is increasing in s. We introduce some notation to simplify the exposition: $f(s) \equiv f(s|V)$, $g(s) \equiv f(s|0)$, and $G(s) = \int_{-\infty}^{s} g(x)dx$.

When uncertain, after observing signal $s \in S$, the examiner forms a belief $Q^*(s)$ that the idea is good. Based on this belief, she chooses a funding probability $\alpha^*(s)$ to solve

$$\max_{\alpha \in [0,1]} \alpha(Q^*(s)(V + c) - c).$$

Thus, $\alpha^*(s) = 1$ if $Q^*(s)(V + c) > c$, $\alpha^*(s) = 0$ if $Q^*(s)(V + c) < c$, and $\alpha^*(s) \in [0, 1]$ if $Q^*(s)(V + c) = c$. In a PBE, beliefs must be consistent with Bayes' rule, so

$$Q^*(s) = \frac{\pi(\sigma, 0)f(s)p}{\pi(\sigma, 0)f(s)p + \pi(\sigma, e^*)g(s)(1 - p)}.$$

By the MLR assumption, $Q^*(s)$ is increasing in s: higher grades increase the posterior belief that the quality of the unrated proposal is good. For any fixed e^*, let us define $\bar{s}(e^*)$ to be the grade such that $Q^*(V + c) = c$, i.e.,

$$\frac{f(\bar{s}(e^*))}{g(\bar{s}(e^*))} = \frac{\pi(\sigma, e^*)(1 - p)c}{\pi(\sigma, 0)pV}. \tag{19}$$

By the MLR assumption,if there exists a solution to this equation it is unique. Then, $\alpha^*(s) = 0$ if $s < \bar{s}(e^*)$, $\alpha^*(s) = 1$ if $s > \bar{s}(e^*)$, and $\alpha^*(s) \in [0, 1]$ if $s = \bar{s}(e^*)$. Thus, we can characterize the examiner's optimal equilibrium strategy by the threshold $\bar{s}(e^*)$.

Conditional on *any* strategy $\alpha(s)$, let $\hat{\alpha} = E_s[\alpha(s)|v = 0]$. The low-type applicant chooses an obfuscation effort e^* to solve

$$\max_{e \geq 0} \hat{\alpha} g_0 \pi(\sigma, e) - e.$$

The solution to this problem is identical to the baseline case. Thus, a perfect-Bayesian equilibrium can be characterized by the expected probability $\hat{\alpha}$ and the obfuscation effort e^*, which determines a threshold $\bar{s}(e^*)$ defined in (19), such that

1. The expected probability of obtaining a patent when the examiner is uncertain is

$$\hat{\alpha} = \int_{\bar{s}(e^*)}^{\infty} g(s)ds = 1 - G(\bar{s}(e^*)).$$

2. The equilibrium obfuscation effort is consistent with $\hat{\alpha}$, so $e^* = e(\hat{\alpha}, \sigma)$. That means that e^* is the solution to the fixed point problem

$$e^* = e(1 - G(\bar{s}(e^*)), \sigma).$$

B.1 Equilibrium Characterization

Recall that in the baseline case there is no obfuscation when $\sigma \leq \bar{\sigma}_1$, which continues to be true in this extension. However, under different degrees of uncertainty obfuscation arises in cases when there was no obfuscation with a single level of uncertainty.

Specifically, let $\sigma > \bar{\sigma}_1$ and $p(V + c) < c$. In the baseline case, there is no obfuscation in this region because the examiner's optimal strategy is $\alpha^* = 0$. Now, since there is a chance that the examiner receives a good signal from a bad idea, when the examiner chooses $\alpha(s) > 0$ for high signals we will have $\hat{\alpha} > 0$, in which case the applicant obfuscates. This will generally be the case unless even with the highest possible grade the sponsor rejects a proposal, i.e., when

$$\frac{f(s)}{g(s)} < \frac{\pi(\sigma, 0)(1 - p)c}{\pi(\sigma, 0)pV}, \quad \text{for all } s \in S,$$

in which case, $e^* = 0$. If this condition above does not hold, and with the lowest signal the examiner rejects, then $\bar{s}(0) < \max\{s \ : \ s \in S\}$, which implies that $\hat{\alpha} > 0$. Then, $e^* > 0$.

Different degrees of uncertainty can also moderate obfuscation. Consider a prior $p > \bar{p}(\sigma)$. In the baseline case, this prior is so high that the examiner grants a patent whenever uncertain. Here, the examiner is optimistic enough to grant a patent for sure even after receiving the lowest possible signal, i.e., when

$$\frac{f(s)}{g(s)} > \frac{\pi(\sigma, e(1, \sigma))(1 - p)c}{\pi(\sigma, 0)pV}, \quad \text{for all } s \in S.$$

Only in this case, $\alpha^*(s) = 1$ for all s and $\hat{\alpha} = 1$, so $e^* = e(1, \sigma)$ is an equilibrium. If this condition does not hold, then $\bar{s}(e(1, \sigma)) < \max\{s \ : \ s \in S\}$, which implies that $\hat{\alpha} < 1$. Then, $e^* < e(1, \sigma)$.

These results show that adding degrees of uncertainty moderates the level of obfuscation in equilibrium: it increases it when $p(V + c) < c$ and it reduces it when $p > \bar{p}(\sigma, g)$.

Finally, consider the case $p \in (c/(V+c), \bar{p}(\sigma))$. In the baseline case, the examiner chooses to grant a patent with probability $\alpha^* \in (0,1)$ so that

$$\frac{\pi(\sigma, e(\alpha^*, \sigma)(1-p)c}{\pi(\sigma, 0)pV} = 1.$$

With different degrees of uncertainty, in equilibrium, (19) must hold .This mean that $e^* = e(\alpha^*, \sigma)$ is the equilibrium effort only if $f(\bar{s}(e^*)) = g(\bar{s}(e^*))$. Otherwise, if $f(\bar{s}(e^*)) > g(\bar{s}(e^*))$ different degrees of uncertainty induces *less* obfuscation than the baseline case, and *more* obfuscation when $f(\bar{s}(e^*)) < g(\bar{s}(e^*))$.

B.2 Comparative Statics

As in the baseline case, e^* will be weakly increasing in σ because taking implicit derivative of the equilibrium condition $e^* = e(1 - G(\bar{s}(e^*)), \sigma)$ we get

$$\frac{de^*}{d\sigma} = -\frac{\partial e(1 - G(\bar{s}(e^*)), \sigma)}{\partial \alpha} g(\bar{s}(e^*)) \frac{\partial \bar{s}}{\partial e^*} \frac{de^*}{d\sigma} + \frac{\partial e(1 - G(\bar{s}(e^*)), \sigma)}{\partial \sigma},$$

which implies

$$\frac{de^*}{d\sigma}\left(1 + \frac{\partial e(1 - G(\bar{s}(e^*)), \sigma)}{\partial \alpha} g(\bar{s}(e^*)) \frac{\partial \bar{s}}{\partial e^*}\right) = \frac{\partial e(1 - G(\bar{s}(e^*)), \sigma)}{\partial \sigma} > 0.$$

Thus, increasing σ increases e^*, since the parenthesis in the LHS is strictly positive. This means that more resources for examiners will generally have an ambiguous impact on the examiner's payoff, since it hinges on the sign of

$$\frac{\partial \pi(\sigma, e^*(\sigma))}{\partial \sigma} + \frac{\partial \pi(\sigma, e^*(\sigma))}{\partial e} \frac{\partial e^*}{\partial \sigma}.$$

C Ex-ante commitment to Patent

The applicant's obfuscation effort is $e(\alpha, \sigma)$, which is increasing in α. Suppose that the examiner *commits* ex-ante to grant a patent with probability α whenever is uncertain.

The optimal choice of α maximizes

$$u_E(\alpha) \equiv p(1 - \pi(\sigma, 0))V + \alpha[p\pi(\sigma, 0)V - (1-p)\pi(\sigma, e(\alpha, \sigma))c]. \tag{20}$$

The term in the bracket corresponds to the expected gain minus the expected loss granting a patent when uncertain: the expected benefit is to patent overlooked good ideas, $p\pi(\sigma, 0)V$, and the expected cost is to patent a bad idea, $(1-p)\pi(\sigma, e(\alpha, \sigma))c$.

First, if $pV \leq c$ the term inside the bracket in (20) is negative because

$$p\pi(\sigma, 0)V \leq (1-p)\pi(\sigma, 0)c \leq (1-p)\pi(\sigma, e(\alpha, \sigma))c$$

since $e(\alpha, \sigma) \geq e(0, \sigma) = 0$ and $\pi(\cdot)$ increases in e. The solution is then to set $\alpha^* = 0$. This is the same probability that the sponsor chooses in the main text without commitment. Thus, there is no gain from commitment when $pV \leq c$.

Second, if $pV > c$, the bracket that multiplies α in (20) is strictly positive when $\alpha = 0$, and decreasing in α. Let

$$\bar{\alpha} = \max\{\alpha \in [0, 1] \ : p\pi(\sigma, 0)V \geq (1-p)\pi(\sigma, e(\alpha, \sigma))c\}.$$

For any $\alpha > \bar{\alpha}$ the bracket is strictly negative and for any $\alpha \leq \bar{\alpha}$ the bracket is positive. Clearly, the value of α^* that maximizes (20) belongs to $[0, \bar{\alpha}]$. Therefore, in general, an examiner with ex-ante commitment grants a patent with a weakly lower probability that she would otherwise in a PBE without commitment, which is $\bar{\alpha}$ (denoted by α^* in the main text). When the solution is strictly lower than $\bar{\alpha}$ there is a strictly positive benefit from ex ante commitment.

Taking the derivative with respect to α we get

$$u_S'(\alpha) = [p\pi(\sigma, 0)V - (1-p)\pi(\sigma, e(\alpha, \sigma))c] - \alpha(1-p)\frac{\partial \pi(\sigma, e(\alpha, \sigma))}{\partial e}\frac{\partial e(\alpha, \sigma)}{\partial \alpha}c.$$

Evaluating this derivative at $\bar{\alpha} < 1$, the term in the bracket is negative and the second term is strictly negative. Therefore, optimal solution is strictly lower than $\bar{\alpha}$. This means, commitment strictly benefit the examiner. If $\bar{\alpha} = 1$, term in the bracket is negative may be positive, which could upset the negative second term, and we could possibly get a corner solution, in which case commitment may not help.

Proposition 10. *Suppose the examiner commits to a probability of approval α_C^* before reading the application. Let $\bar{\alpha}$ be the equilibrium probability without commitment (as in the main text). Then, we have: (1) when $pV \leq c$, then $\alpha_C^* = \bar{\alpha} = 0$; (2) when $pV > c$ and $\bar{\alpha} < 1$, then $\alpha_C^* < \bar{\alpha}$; (3) when $pV > c$ and $\bar{\alpha} = 1$, $\alpha_C^* \leq \bar{\alpha}$. Thus, with ex-ante commitment, the examiner rejects (weakly) more when $pV > c$.*

Next, we study the impact of commitment on the marginal benefit of better referees. Intuitively, with commitment, the sponsor should be better off with better referees. By the

envelop theorem, the derivative of the sponsor with respect to σ is simply

$$\begin{aligned}\frac{du_E(\sigma,\alpha^*(\sigma))}{d\sigma} &= \frac{\partial u_E(\sigma,\alpha^*)}{\partial\sigma} \\ &= -(1-\alpha^*)p\pi_\sigma(\sigma,0)V - \alpha^*(1-p)c\left[\pi_\sigma(\sigma,e(\alpha^*,\sigma)) + \pi_e(\sigma,e(\alpha^*,\sigma))\frac{de(\alpha^*,\sigma)}{d\sigma}\right].\end{aligned}$$

In the main text, whenever the term in the bracket above was positive, additional resources reduced the sponsor's payoff. Here, the first term counteracts the examiner's potential loss from additional resources caused by the increase in obfuscation.

Corollary 1. *Relative to the case of no commitment, the benefit of more resources is weakly larger under commitment.*

D Supermodularity

In this section, we study how our results would change if we, instead, assume that $\pi(\sigma,e)$ submodular. That is, the marginal return to effort is *lower* when the applicant faces a more resourceful examiner. Again, only the low-type applicant has incentives to obfuscate and chooses her obfuscation effort by solving

$$\max_{e\geq 0}\ \alpha\cdot g_0\cdot\pi(\sigma,e) - e.$$

In contrast to the baseline case, the solution to this problem, $e(\alpha,\sigma)$, is weakly *decreasing* in σ. First, if $e(\alpha,\sigma)=0$, then $e^*(\alpha,\sigma')=0$ for any $\sigma'>\sigma$.

In general, there is a threshold $\bar{\sigma}_1$ (which could be $\sigma=1$) such that for any $\sigma\geq\bar{\sigma}_1$, we have $e(1,\sigma)=0$. Thus, when $\sigma\geq\bar{\sigma}_1$, there is no obfuscation, and when $\sigma<\bar{\sigma}_1$, there is obfuscation when $\alpha=1$.

In contrast to the previous case, suppose that $\alpha=1$. Note that $\pi(\sigma,e(1,\sigma))$ is *decreasing* in σ, because $e(1,\sigma)$ decreases and the direct effect, π_σ, is negative. That means Q^* is increasing in σ.

As in the baseline case, $Q^*(V+c)>c$ if and only if

$$\pi(\sigma,0)pV > \pi(\sigma,e^*)(1-p)c. \tag{21}$$

When $e^*=0$, this condition is equivalent to $pV>c$. Thus, for any $\sigma\geq\bar{\sigma}_1$ this condition holds, so $\alpha^*=0$, $e^*=0$, and $Q^*=p$ is an equilibrium.

Note that the slope of the left-hand side in (21) (as a function of σ) is less negative that the slope of the right-hand side (because $e(\alpha, \sigma)$ decreases in σ). When $\pi(0,0)pV \geq \pi(0, e(1,0))(1-p)c$, $\alpha^* = 1$ always. When $\pi(0,0)pV < \pi(0, e(1,0))(1-p)c$, if $pV > c$, in equilibrium $\alpha^* = 1$ and $e^* = e(1, \sigma)$ for $\sigma \in [\bar{\sigma}_2, \bar{\sigma}_1]$, for some threshold $\bar{\sigma}_2$. If $pV > c$ and $\sigma < \bar{\sigma}_2$, in equilibrium α^* to solves

$$\pi(\sigma, 0)pV = \pi(\sigma, e(\alpha^*, \sigma))(1-p)c.$$

Finally, we have the case $pV < c$ and $\sigma < \bar{\sigma}_1$. In this case, $\alpha^* = 0$. Figure A.3 shows that the taxonomy of equilibria is similar to the previous case. However, the comparative statics results reverse. The examiner's equilibrium payoff can also be analyzed in two cases, $\alpha^* = 1$ or $\alpha^* < 1$, and it is now increasing everywhere because

$$\underbrace{\frac{\partial \pi(\sigma, e^*)}{\partial \sigma}}_{<0} + \underbrace{\frac{\partial \pi(\sigma, e^*)}{\partial e}}_{>0} \underbrace{\frac{\partial e^*}{\partial \sigma}}_{<0} > 0.$$

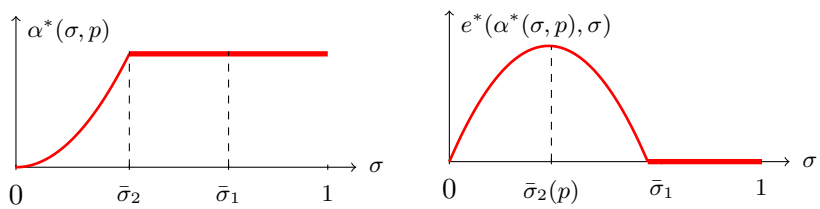

Figure A.3: The figure considers the case $pV > c$. The left plot shows the equilibrium probability of receiving a patent from an uncertain examiner as a function of the examiner's resources. The right plot shows the equilibrium obfuscation effort $e(\alpha^*, \sigma)$ as a function of σ.

Figure A.4 shows that the difference in equilibrium payoffs for applicants of different types as a function of p. The figure shows that the payoff difference has the same features as the baseline case. This implies that the equilibrium when applicants search for ideas will exist and it will be unique. The only difference is that the regions are determined by different values of σ, which will now change the equilibrium comparative statics results.

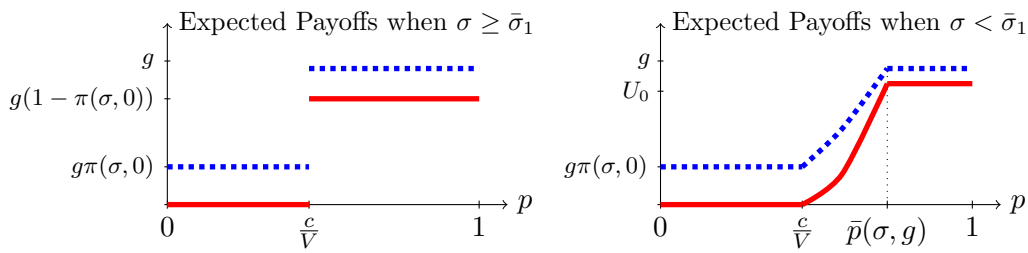

Figure A.4: The figure consider the case $pV > c$ and shows the applicants' equilibrium expected payoffs as a function of p. The left plot is for $\sigma \geq \bar{\sigma}_1$ and the right plot for $\sigma < \bar{\sigma}_1$. The dashed line is payoff for the applicant with a good idea (U_V) and the solid line is payoff for the applicant with a bad idea (U_0). We denote by $U_0 = (1 - \pi(\sigma, e(1,\sigma)))g - e(1,\sigma)$.

E Heterogenous Private Values

Let us consider the case $g_V > g_0$. The main difference with the main text is the value of different ideas:

$$U_0(p) = \alpha^* \pi(\sigma, e^*) g_0 - e^*, \tag{22}$$

$$U_V(p) = [1 - \pi(\sigma, 0) + \alpha^* \pi(\sigma, 0)] g_V. \tag{23}$$

In this case, we have

Proposition 11. *The difference $U_V - U_0$ is weakly decreasing in p whenever*

$$g_V \pi(\sigma, 0) < g_0 \pi(\sigma, e^*),$$

where $e^* = e(\alpha^*(p), \sigma)$.

Proof. We have $U_V(p) - U_0(p) = g_V(1 - \pi(\sigma, 0)) + g_V \alpha^* \pi(\sigma, 0) - g_0 \alpha^* \pi(\sigma, e^*) + e^*$.Taking the derivative with respect to p and using the envelope theorem we get

$$(U_V - U_0)'(p) = \frac{\partial \alpha^*}{\partial p}[g_V \pi(\sigma, 0) - g_0 \pi(\sigma, e^*)].$$

□

The difference $g_V \pi(\sigma, 0) - g_0 \pi(\sigma, e^*)$ is the value of the benefit of the doubt for an agent with a good idea relative to that value for an agent with a bad idea. In the main text (Section 4), we assume that $g_V = g_0$, in which case the condition in the proposition holds since

$\pi(\sigma, 0) < \pi(\sigma, e^*)$. This assumption reflects that the benefit of the doubt is more valuable for agents with bad ideas.

When $g_V > g_0$ the benefit of the doubt is more valuable for agents with good ideas for some values of p. That is, $U_V(p) - U_0(p)$ is generally non-monotone. If g_V is very large relative to g_0, $U_V(p) - U_0(p)$ is increasing. However, if g_V is not too large relative to g_0, $U_V(p) - U_0(p)$ will be weakly decreasing in p except when p is in a small interval $p \in [c/(V + c), p^n]$.

Even when $U_V(p) - U_0(p)$ is non-monotone, equilibrium existence is guaranteed. Uniqueness holds under an additional parametric condition, characterized in the proposition below.

Proposition 12. *There exists an equilibrium (p^*, k^*). The equilibrium is unique when*

$$\mu_0 + F((\mu_S - \mu_0)[U_V(\sigma, p, g) - U_0(\sigma, p, g)])(\mu_S - \mu_0) - p$$

satisfies single crossing.

Proof. An equilibrium (p, k) solves:

$$\begin{aligned} p &= \mu_0 + F(k)(\mu_S - \mu_0) \\ k &= (\mu_S - \mu_0)[U_V(p) - U_0(p)] \end{aligned}$$

We define $H(p) = \mu_0 + F((\mu_S - \mu_0)[U_V(\sigma, p, g) - U_0(\sigma, p, g)])(\mu_S - \mu_0)$. In equilibrium, p solves the fixed-point equation $p = H(p)$. $H(\cdot)$ is a continuous and maps $[0, 1]$ into $[0, 1]$. By the intermediate value theorem, an equilibrium exists. If $H(p) - p$ satisfies single crossing, then clearly there is a unique equilibrium. □

Our results hinge on the condition $g_V \pi(\sigma, 0) < g_0 \pi(\sigma, e^*)$ holding at the equilibrium. This requires that $U_V(p^*) - U_0(p^*)$ decreases at the equilibrium. A few conditions guarantee that this will be the case. For instance, whenever $|g_V - g_0| < \ell$, with ℓ sufficiently small, $g_V \pi(\sigma, 0) < g_0 \pi(\sigma, e(1, \sigma))$, and $c/(V + c) < \kappa$, the equilibrium will occur in the interior of the region of parameters where $U_V - U_0$ decreases. Hence, the results of the main text continue to hold in this setting, albeit with further qualifications.