WORKING PAPERS

N° TSE-442

August 2013

"Fear of Rejection? Tiered Certification and Transparency"

Emmanuel Farhi, Josh Lerner and Jean Tirole

Institut d'Economie Industrielle
Toulouse School of Economics

Fear of Rejection?*

Tiered Certification and Transparency

Emmanuel Farhi Josh Lerner Jean Tirole

August 27, 2013

Abstract

Product quality certifiers may not reveal the identity of unsuccessful applicants/sellers for three reasons. First, they respond to the desire of individual sellers to avoid the stigma from rejection. Second, non-transparency helps a certifier to increase his market power by raising the stigma from lower-tier certification. Third, transparency does not help screen among heterogeneous sellers. Strategic complementarities arise as sellers move down the certification pecking order and lead to the stigmatization of the lower tiers. Mandating transparency benefits the sellers, but has an ambiguous impact on buyers, who actually become less informed about product quality.

Keywords: certification industry, transparency, stigma, rejections.

JEL numbers: D82, 031, 034.

*Respectively Harvard University (efarhi@fas.harvard.edu), Harvard University (jlerner@hbs.edu) and Toulouse School of Economics (TSE) and Institute for Advanced Study in Toulouse (IAST) (jean.tirole@tse-fr.eu). We thank several seminar audiences, two anonymous referees, and Judith Chevalier for helpful comments. The U.S. National Science Foundation (NSF Grant #0830288, "Patent Pools and Biomedical Innovation") and Harvard Business School's Division of Support provided financial support. Emmanuel Farhi and Jean Tirole acknowledge funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013) Grant Agreement no. 249429. All errors and omissions are our own.

1 Introduction

As most markets are characterized by imperfect knowledge, informational intermediaries have become central to their working. From underwriters to rating agencies, from scientific journals to entry-level examinations, from standard-setting organizations to system integrators, intermediaries serve sellers and buyers by providing product-quality information to the latter.

The literature on intermediaries has carefully analyzed their incentives to collect product-quality information and disclose it truthfully. By contrast, we know little about three related aspects of the certification market: the publicity given to applications (i.e., the transparency or opacity regarding rejections), the coarseness of rating patterns, and the sellers' sequential certification strategies. Policies in these matters exhibit substantial heterogeneity.

Regarding the *transparency* of the application process, scientific journals, certified bond rating agencies, lenders, underwriters, employers, universities, and organic food certifiers usually do not reveal rejected applications. By contrast, some entry-level examinations (SAT, GMAT,...) have historically disclosed previous, and presumably unsuccessful, attempts by the student. In 2009, though, the College Board began allowing students to only report certain SAT test scores to colleges, rather than all results as previously. Critics questioned whether this program was a competitive response to the competing ACT test, which has long had a similar policy of non-transparency. For instance, Stanford's admissions head asked "Was this a student-centered decision? Or was this business-centered because they're worried about losing market share?"[1]

Regarding the *coarseness of grading*, many institutions, such as most scientific journals, adopt a "minimum standard" or "pass-fail" strategy, while others, such

[1] See Rimer (2008).

as entry-level examination firms, report an exact grade. While a fine partition in the grading space presumably requires more resources than a pass-fail approach, what drives the choice of coarseness is unclear.

While we will consider a variety of examples, the illustration that we will return to repeatedly is the certification of papers by academic journals. We here focus on the role of journals as certifiers (informational intermediaries) more than as distributors. While the two functions have historically been bundled, they need not be. The relative importance of certification has increased with easy access to alternative distribution modes, particularly in fields such as economics which allow free circulation of unpublished working papers.[2] The publication process is opaque as almost all journals refrain from publishing the list of submitted or rejected papers; they also by and large follow a minimum standard approach, even though a lead article carries some added prestige.

But our analysis applies to a number of other industries. Table 1 reports the strategies of some certifiers regarding publicity and grading. "Application opacity" refers to the certifier's policy of not disclosing rejected applications, not necessarily to the outcome.[3] Consider, for example, the recent efforts in the United States to ensure transparency of the securities rating process, particularly in the area of structured finance. Issuers have in a number of ways been able to get rating

[2] Our key insights hold whether the certifier charges buyers or sellers, and whether they maximize profit or pursue a non-profit goal such as market share.

[3] For example, one may fortuitously learn that a paper was submitted to and rejected by a journal; furthermore, a delayed publication may create some stigma as the profession is unsure as to whether the delay is due to the author, slow editing or a rejection. This is likely to be less of an issue in many journals in the hard sciences, where working papers are often embargoed prior to publication. Similarly, while academic departments, corporations and partnerships warn in advance assistant professors and junior members that they are unlikely to receive tenure or keep their job, thereby allowing them to attempt to disguise a layoff as a resignation, information leakages and the inference drawn from the very act of quitting limit this strategy. Relatedly, faculty members who want to quit but expect to be successful at a particular promotion or reappointment will often delay their resignation until after their review process is completed, as they want to avoid the negative inference associated with an early resignation.

agencies to not disclose adverse ratings. First, even if a firm appeals a rating that displeases it and the appeal is rejected, the proposed rating still may not be published. Instead, a "break-up fee" is paid by the issuer to the rating agency to compensate it for its efforts (U.S. Securities and Exchange Commission, 2008). Alternatively, consulting services offered in recent years by rating agencies to issuers may have made an apparently transparent process opaque.[4]

TABLE

As for *sequential certification strategies*, sellers often adopt an ambitious or top-down submission strategy, in which they apply first to demanding and non-transparent certifiers and then, in case of rejection, move down the pecking order.[5] Why do we observe this pattern, and what determines the rejection rate, or equivalently whether submissions tend to be ambitious or realistic?

This paper develops a framework in which the descriptive and public policy aspects of transparency and grading strategies can be initially analyzed. Needless to say, our model does not capture all elements of each of the many relevant situations, but only some key trade-offs that they share. It builds on the idea that certifiers' policies must reflect the demands of the two sides of the market, as well as who has "gatekeeping power" over the certification process. In the majority of applications in Table 1, the seller chooses the certifier.[6] While they need to be

[4] "*The inherent conflict facing the credit rating agency has been aggravated by their recent marketing of advisory and consulting services to their clients. Today, the rating agencies receives one fee to consult with a client, explain its model, and indicate the likely outcome of the rating process; then, it receives a second fee to actually deliver the rating (if the client wishes to go forward once it has learned the likely outcome). The result is that the client can decide not to seek the rating if it learns that it would be less favorable than it desires; the result is a loss of transparency to the market.*" (Coffee 2008). See also Partnoy (2006).

[5] An exception to this widespread pattern is provided by publications in law journals, where authors build on acceptance to move up the quality ladder.

[6] Interestingly in view of the pattern exhibited in Table 1, the seller does not choose the certifier in the case of state licensing and professional exams.

credible vis-à-vis the buyers, the certifiers must first cater to the sellers' desires.

At an abstract level, a certifier's policy maps the information it acquires about the quality of the product into a public signal; if specified in the contract with the seller, the certifier can refrain from providing a signal, thereby concealing the existence of an application in order not to convey bad news about quality. We also allow for fortuitous disclosure, as buyers may hear about the application "through the grapevine" even though the certifier does not disclose it.

We find conditions under which sellers opt for an ambitious strategy. We derive the comparative statics with respect to the sellers' initial reputation, the probability of fortuitous disclosure, the sellers' self-knowledge and impatience, and the concentration of the certification industry. Finally, we investigate the implications of regulating transparency.

The paper's key insights are summarized in the conclusion, where we also revisit the patterns in Table 1. A key theme of the paper is that sellers have a clear *individual* preference for non-transparency, as they do not want buyers to know about their rejections. Tier-1 certifiers with market power have, as we will show, two other motives for preferring non-transparency. A second broad theme is that certification strategies exhibit strategic complementarities through the stigma associated with second-tier certification: When certification by second-tier institutions carries a big stigma, sellers are tempted to first aim high, thereby confirming the buyers' concern that second-tier certifiers attract previously rejected applications. A third theme relates to public policy. Non-transparency collectively always hurts sellers (while its impact on buyers is ambiguous).

The paper is organized as follows. Section 2 lays out the basic model, in which only minimum-standard certification is offered. Section 3 solves for the equilibrium in a competitive or concentrated certification industry and conducts the welfare analysis of transparency regulation. Section 4 develops a number of extensions: It analyzes the impact of seller heterogeneity and allows for multi-tier grading. Section 5 summarizes the insights and discusses a number of open questions. Finally, the Appendix considers three additional extensions: It generalizes the basic model by endogenizing the sellers' quality choice; it allows certifiers to charge buyers; and finally it examines the effect of entry by certifiers who trade off accuracy and turn-around time.

Relationship to the literature

There is a large literature on certification in corporate finance, industrial organization and labor markets. Much of this literature focuses on the trade-off for parties seeking certification between the cost of certification and its benefits in terms of signaling, reduced agency costs or assortative matching. Much less has been written on the industrial organization of the certifying industry.

A series of papers have studied the coarseness issue, without linking to opaqueness. Lizzeri's (1999) classic paper analyzes the choice of coarseness in an environment in which the seller perfectly knows the quality of her product and applications to a certifier are (endogenously) public. Lizzeri shows that a monopoly certifier may either disclose nothing at all or go for a minimal standard, while a competitive certifying industry can lead to full disclosure of quality. Faure-Grimaud et al. (2009), in a model in which sellers are imperfectly informed about the quality of their product, allow sellers to "own ratings," in the sense that they can hide the rating from the buyers if they choose to. They show that the imprecision of the rating technology per se does not lead to certifiers offering the concealment option to sellers. Rating ownership arises only if sellers are very unsure about their product's quality. In Lerner–Tirole (2006), certifiers differentiate through their composition and decision-making processes, making them more or less friendly to sponsors' interests. The current paper investigates certifiers' positioning with respect to transparency; it further analyzes sequential rejections, an issue that does not arise when the technology sponsor's objective is simply to have the technology adopted, as in Lerner-Tirole.

In environments such as those considered in Lizzeri and Lerner-Tirole, in which either the seller has a single chance to be certified or the seller's application is public, the certifier cannot or does not want to conceal applications. Faure-Grimaud et al. do allow for the possibility of hiding the rating, but do not consider the possibility of a second chance, perhaps by going to a less demanding certifier.

Morrison and White (2005) and Gill and Sgroi (2004) allow for a second chance, but do not focus on the opacity issue. Banks in Morrison-White apply to regulators with different perceived abilities. A successful application to a tough regulator allows banks to raise more deposits. As regulators make mistakes, banks may get a second chance. On the other hand, the Morrison-White paper focuses on rather different issues than our paper; for instance, it assumes that applications are transparent. By contrast, applications are non-transparent in Skreta-Veldkamp's (2009) work on rating ownership; in that paper, sellers shop for ratings and disclose them selectively. Each rating agency issues an unbiased forecast of the product quality (there is no coarseness choice). Buyers are naïve. The selective disclosure of ratings by sellers matters more, the more mistakes certifiers make; Skreta and Veldkamp thereby provide a narrative on the recent treatment of complex assets by financial rating agencies. In a related vein, Sangiorgi et al. (2009) analyze the winner's curse problem associated with more sophisticated buyers and the selective disclosure of ratings. Bolton-Santos-Scheinkman (2011) analyzes interactions between an opaque OTC market and a transparent exchange, in which the OTC market may siphon out good assets and lower the perceived quality of assets in the exchange, leading to a multiplicity of equilibria similar to that unveiled in Proposition 1 below. Bolton et al. study themes, such as dealer rents and the talent allocation to finance, that we do not pursue.

Like these papers, we assume that certifiers have sufficient governance or reputation to abide by the reporting rule that they announce.[7] Certifiers' reputation building is analyzed in Biglaiser (1993), Bolton-Freixas-Shapiro (2011), Bouvard-Levy (2009) and Mathis et al (2009). Relatedly, our focus on coarseness relates to the very extensive literature on cheap talk. But unlike in the cheap talk literature, there is here no issue of credibility of the reporting strategy, as certifiers are as truthful as their reporting rule allows them to be. Coarseness thus stems from the sellers' request, not from moral hazard and an a priori misalignment between certifier and buyers.

2 The model

2.1 Description

Time is discrete and runs from $-\infty$ to $+\infty$. There is a mass 1 and a steady inflow of sellers, each with one product of unknown single-dimensional quality. For simplicity, the representative seller's quality i, which is initially unknown to both sides of the market, can take one of three values: high (H), low (L) or "abysmal" ($-\infty$), with respective benefits for the buyers $b_i \in \{b_H, b_L, -\infty\}$ with

[7] Of course, this assumption does not always hold in the real world. For instance, some critics have accused rating agencies of initially being excessively generous when rating new offerings, then revising the rating months later. They suggest that the natural organizations to question this behavior, the investment banks, had little incentive to do so, because they had typically "laid off" any exposure to the securities through refinancings (U.S. Securities and Exchange Commission, 2003).

$b_H > b_L > -\infty$.[8] Conditional on not being abysmal, quality is high with prior probability ρ and low with prior probability $1 - \rho$. Buyers prefer quality H to quality L, and do not consider the product unless its quality has been certified to be at least L. A seller whose quality has not been certified to be at least L does not bring the product to the market and obtains zero profit.

Assuming that this certification has taken place, let $\widehat{\rho}$ denote the buyers' posterior belief at the time at which the product is brought to the market (more on this shortly). Let $S_i(\widehat{\rho})$ denote the seller's expected gain from putting a product of quality i on the market when beliefs are $\widehat{\rho}$. We will assume that S_i is always positive and is weakly increasing in $\widehat{\rho}$.

Certifiers. Profit-maximizing certifiers audit quality. Throughout the paper, we will assume that, whether through reputation or a credible internal-audit mechanism, certifiers are able to commit to a disclosure policy, that is to a mapping from what they learn to what they disclose to buyers. This ability to commit to a disclosure policy makes the question of the structure of their incentive scheme moot,[9] and so we can assume without loss of generality that they demand a fixed fee for the certification service. To sum up, a certifier's strategy is the combination of a fixed fee and a disclosure policy. (Our results are qualitatively the same if we alternatively assume that certifiers do not charge fixed fees and that their objective is to maximize market share. When certifiers are atomistic and competition is perfect, the outcome is exactly the same. Differences potentially materialize when we consider monopolistic competition.)

Because certifiers are useless unless they rule out the abysmal quality, we can consider three types of certifiers, two "minimum standard" certifiers and one "full

[8] The role of the assumption that quality can be $-\infty$ is to ensure that goods must be certified. We could alternatively have directly made the assumption that goods must be certified.

[9] An arbitrary incentive scheme can be duplicated through a fixed payment (equal to the expected payment under the incentive scheme) and the same commitment to a disclosure policy.

grade" or "multi-tier" certifier:

A *tier-1* certifier ascertains that $b = b_H$ or $b \in \{b_L, -\infty\}$. Tier-1 certifiers furthermore do not disclose applications for which they find that $b \in \{b_L, -\infty\}$, as such disclosure of bad news (a "rejection") is unappealing to sellers and reduces the demand for such certifiers' services.

A *tier-2* certifier certifies that $b \in \{b_H, b_L\}$ or $b = -\infty$.[10]

A *multi-tier* certifier discloses the true quality: $b = b_H, b_L$ or $-\infty$.

We will normalize the audit cost incurred by a minimum standard certifier to be 0. By contrast, the cost of a finer grading may be positive (see Section 4.3). Certifiers compete for the sellers' business. The certification market, unless otherwise stated, is perfectly competitive. Equilibrium fees are then equal to 0.

Timing. Consider a seller who arrives at date t and chooses a certifier. She can contract with a single certifier in each period. Contingent on the outcome of certification(s), the seller chooses the date, $t + \tau$ ($\tau \geq 0$), at which she brings the product to the market. If the buyers' beliefs at that date are $\widehat{\rho} = \widehat{\rho}_{t+\tau}$, then the seller's utility is

$$\delta^\tau S_i(\widehat{\rho}_{t+\tau})$$

where $\delta < 1$ is the discount factor. Thus the seller maximizes

$$E[\delta^\tau S_i(\widehat{\rho}_{t+\tau})].$$

In our model, there are only two (relevant) levels of quality and audits of a given

[10] Obviously, the certifier's reporting strategy for $b = -\infty$ is irrelevant, as the seller then always makes no profit. If, by contrast, we assumed that sellers have other products, the production of an "abysmal quality" could be a bad signal for other offerings. One would then expect that the information that $b = -\infty$ would not be disclosed either.

kind always deliver the same outcome.[11] And so a date-t product will actually be brought to the market either at t or at $t+1$.

There can be fortuitous disclosure: When a seller arrives at date t and does not bring her product to the market until date $t+1$, with probability $d \geq 0$, buyers exogenously discover that the date-$(t+1)$ introduction corresponds to a date-t arrival. With probability $1-d$, buyers receive no such information.[12]

Equilibrium: We will analyze perfect Bayesian equilibria. If multiple equilibria coexist, which can be Pareto ranked for the sellers, we will select the Pareto dominant one.

Sellers' attitude toward information revelation. In general, sellers may, ceteris paribus, welcome or dread the revelation of information about quality. Our key insights do not depend on which obtains, but some comparative statics results can be obtained with regard to seller attitudes in this respect. We will say that the seller is information loving (respectively, averse) if in a situation in which she could not get a second chance, she would prefer full revelation (respectively, a coarse disclosure specifying that quality is at least L). Thus, sellers are *information loving* if $\rho S_H(1) + (1-\rho)S_L(0) > \rho S_H(\rho) + (1-\rho)S_L(\rho)$; *information averse* if $\rho S_H(1)+(1-\rho)S_L(0) < \rho S_H(\rho)+(1-\rho)S_L(\rho)$; *information neutral* if $\rho S_H(1)+(1-\rho)S_L(0) = \rho S_H(\rho)+(1-\rho)S_L(\rho)$. Note that when $S_i(\rho) = S(\rho)$ is independent of i, sellers are information loving if and only if $S(\rho)$ is strictly convex, information averse if and only if $S(\rho)$ is strictly concave, and information neutral if and only if

[11] There is no certifier-idiosyncratic noise, unlike in Morrison-White (2005) or Skreta-Veldkamp (2009).

[12] Fortuitous disclosures will in equilibrium decrease the seller's utility when being rejected. Note that learning that the seller arrived at date t is here equivalent to learning that her application was rejected at date t. We could easily enrich the model by adding "slow sellers," who arrive at date t, but apply only at date $t+1$. Such sellers would suffer an unfair stigma if the date of their arrival is made public, as do papers in academia that authors are slow at submitting to a journal.

$S(\rho)$ is affine. Section 2.2 provides some illustrations of preferences with respect to information. For example, a demand for assortative matching between buyers and sellers tends to make sellers information-loving. Belief risk aversion or the presence of complementary sales push in the direction of information aversion.

2.2 Examples

Example 1 (sale). Suppose that production is costless and that the seller sells the product to homogenous, price-taking consumers. Then, under such first-degree price discrimination

$$S_i(\widehat{\rho}) = \max\{E_{\widehat{\rho}}[b], 0\}$$

is independent of i, where $E_{\widehat{\rho}}[b] \equiv \widehat{\rho} b_H + (1-\widehat{\rho}) b_L$ denotes the consumers' posterior assessment of quality.

Example 2 (sale with imperfect price discrimination). Following up on Example 1, assume now that there are two types of consumers, indexed by $a = a_H$ (proportion μ) or a_L (proportion $1-\mu$) with $a_H > a_L$. If $\widehat{b} = E_{\widehat{\rho}}[b]$, the gross surplus of a user of type $j \in \{H, L\}$ is $a_j + \widehat{b}$. "Belief-sensitive pricing" (which would be the "generic case" with a continuum of types) arises when user surplus depends on posterior beliefs $\widehat{\rho}$,[13] i.e., when

$$a_L + b_H > \mu(a_H + b_H) \text{ and } a_L + b_L < \mu(a_H + b_L).$$

[13] The other two cases are isomorphic to Example 1, as the volume of sales is not affected by beliefs.

Then, $S_i(\widehat{\rho})$ (which again is independent of i) is given by

$$S_i(\widehat{\rho}) = \begin{cases} a_L + \widehat{b} & \text{for } \widehat{\rho} \geq \rho_0 \\ \mu(a_H + \widehat{b}) & \text{for } \widehat{\rho} < \rho_0 \end{cases}$$

where

$$a_L + [\rho_0 b_H + (1-\rho_0) b_L] = \mu[a_H + \rho_0 b_H + (1-\rho_0) b_L].$$

Buyers then have (average) utility

$$B(\widehat{\rho}) = \begin{cases} \mu(a_H - a_L) & \text{for } \widehat{\rho} \geq \rho_0 \\ 0 & \text{for } \widehat{\rho} < \rho_0 \end{cases}.$$

Example 3 (clientele effects / assortative matching). Some buyers may have a strong preference for high-quality offerings due to prudential regulation. For example, many public pension funds are allowed to hold only investment-grade securities. Full grading allows the seller to better segment the market. Suppose that a fraction of buyers buy only high-quality products, at price Kb_H where $K > 1$. Other buyers are less discriminating and are as depicted in Example 1. Then

$$S_i(\widehat{\rho}) = Kb_H \mathbb{1}_{\{\widehat{\rho}=1\}} + \max\{E_{\widehat{\rho}}[b], 0\} \mathbb{1}_{\{\widehat{\rho}<1\}},$$

is again independent of i.

Example 4 (spillovers from adoption). A researcher whose paper is read and used by the profession, or a technology sponsor whose intellectual property becomes part of a royalty-free standard benefit only indirectly from adoption (prestige, referencing, diffusion of ideas for a researcher, network effects or spillovers to complementary products for a technology sponsor). Letting s_i denote the seller's gross benefit from adoption. The seller's surplus is then:[14]

$$S_i(\widehat{\rho}) = s_i \mathbb{1}_{\{E_{\widehat{\rho}}[b] \geq 0\}}.$$

Note that in this case the seller's surplus in general depends directly on quality i.

If $b_L \geq 0$, the seller is information neutral in Examples 1 and 4, and information loving in Example 3. If $b_L < 0$, she is information loving when she fully appropriates the consumer surplus through a price (Examples 1 and 3). Similarly, the seller is information-loving or neutral in Example 2.

By contrast, the seller is information averse if $E_\rho[b] > 0$ and if she is unable to charge the buyer and therefore has buyer adoption as her primary objective. The seller then always benefits from a no grading, simple-acceptance policy (see Lerner-Tirole 2006), weakly so in the two-type case when $b_L \geq 0$ (as in Example 4) and strictly so with two types and $b_L < 0$ or with a continuum of types, some of them negative. That way, she is able to "pool" negative-buyer-surplus states with positive-buyer-surplus ones.

3 Minimum standard certifiers

3.1 Equilibrium behavior with competitive certifiers

There is no point for a seller to apply to a tier-2 certifier unless she goes to the market following an endorsement. Similarly, after an application to a tier-1 certifier, the seller brings the product to the market if the latter is a high-quality one and applies to a tier-2 certifier in case of rejection. The equilibrium thus exhibits the familiar pattern of moving down the pecking order, with diminishing

[14] Where $\mathbb{1}_{\{\cdot\}}$ is the indicator function.

expectations.

Let x denote the fraction of sellers who choose an *ambitious strategy* (start with a tier-1 certifier, and apply to a tier-2 certifier in case of rejection). Fraction $1-x$ select the *safe strategy* (go directly to a tier-2 certifier).

When faced with a product certified by a tier-2 certifier, buyers form beliefs:

$$\begin{cases} \widehat{\rho} = 0 & \text{if they know the product introduction is delayed (as they then infer} \\ & \text{a rejection in the previous period), and} \\ \widehat{\rho} = \widehat{\rho}(x) \equiv (1-x)\rho/\left[1-x+x(1-\rho)(1-d)\right] & \text{otherwise.} \end{cases}$$

Note that $\widehat{\rho}(x)$ decreases from ρ to 0 as x increases from 0 to 1.

Let

$$W^1(\widehat{\rho}) \equiv \rho S_H(1) + (1-\rho)\delta[dS_L(0) + (1-d)S_L(\widehat{\rho})]$$

and

$$W^2(\widehat{\rho}) \equiv \rho S_H(\widehat{\rho}) + (1-\rho)S_L(\widehat{\rho})$$

denote the expected payoffs[15] when applying to a tier-1 or tier-2 certifier, when certification by a tier-2 certifier delivers reputation $\widehat{\rho}$. A crucial property is that, due to (a) discounting, (b) fortuitous disclosure, and (c) the fact that high types pass tier-1 certification, the payoff from tier-2 certification is more sensitive to the perception of quality following tier-2 certification than the payoff from tier-1 certification:

$$\frac{\partial W^2}{\partial \widehat{\rho}} \geq \frac{\partial W^1}{\partial \widehat{\rho}} \geq 0.$$

- *Safe-strategy equilibrium.* It is an equilibrium for sellers to all adopt a safe

[15] Conditional on $b \in \{b_L, b_H\}$.

strategy $(x = 0)$[16] if $W^2(\rho) \geq W^1(\rho)$:

$$\rho S_H(\rho) + (1-\rho)S_L(\rho) \geq \rho S_H(1) + \delta(1-\rho)[(1-d)S_L(\rho) + dS_L(0)],$$

or

$$(1-\rho)[(1-\delta)S_L(\rho) + \delta d[S_L(\rho) - S_L(0)]] \geq \rho[S_H(1) - S_H(\rho)]. \tag{1}$$

Condition (1) captures the costs and benefits of a safe strategy. A safe strategy avoids delaying introduction when quality is low, thereby economizing $(1-\delta)S_L(\rho)$. It also prevents the stigma associated with fortuitous disclosure, and thereby provides gain $\delta d[S_L(\rho) - S_L(0)]$. The cost of a safe strategy is of course the lack of recognition of a high quality and the concomitant loss $S_H(1) - S_H(\rho)$.

Unsurprisingly, a safe-strategy equilibrium is more likely to emerge, the lower the discount factor (e.g., the longer the certification length), and the higher the rate of fortuitous disclosure. Indeed, when $\delta = 1$, the safe-strategy equilibrium never exists (i.e., even for $d = 1$) if the seller is information loving.

- *Ambitious-strategy equilibrium.* Next, consider an equilibrium in which all sellers adopt an ambitious strategy. Certification by a second-tier certifier is then very bad news. Thus $x = 1$ is an equilibrium if and only if $W^1(0) \geq W^2(0)$:

$$\rho S_H(1) + \delta(1-\rho)S_L(0) \geq \rho S_H(0) + (1-\rho)S_L(0) \tag{2}$$

- *Mixed-strategy equilibrium.* Finally, consider a mixed equilibrium in which

[16] Observed certifier rankings always start with "tier-1", almost by definition. One interpretation of our model is that this particular class of sellers (products) is an identifiable subgroup of sellers (products), who in equilibrium apply to tier-2 certifiers (on this, see also Section 4.2 below). Another interpretation speaks to the very definition of "tier-1", "tier-2", etc. What we here call "tier-2" could in practice be called "tier-1" if no seller applied to what we define as "tier-1" certifiers. For example, no "super tier-1" journal has been created that would be more demanding than the top-5 economics journals and only take, say, the ten best papers of the year.

$x > 0$ (some sellers adopt an ambitious strategy), that is $W^1(\widehat{\rho}(x)) = W^2(\widehat{\rho}(x))$:

$$\rho S_H(1) + \delta(1-\rho)[(1-d)]S_L(\widehat{\rho}(x)) + dS_L(0)] = \rho S_H(\widehat{\rho}(x)) + (1-\rho)S_L(\widehat{\rho}(x)). \quad (3)$$

Condition (3) has at most one solution x. Note also that whenever a mixed equilibrium exists, the safe-strategy equilibrium also exists, and it dominates the mixed equilibrium from the point of view of the sellers. In both a mixed equilibrium and a safe-strategy one, the safe strategy is an optimal strategy; but the stigma is weaker in a safe-strategy one and so sellers are better-off under the safe-strategy equilibrium. The seller-centric Pareto-dominance criterion implies that we can ignore mixed equilibria from now on.

Note that an individual seller does not want to apply to a tier-1 certifier who discloses rejections, i.e., a transparent tier-1 certifier. If the equilibrium features ambitious strategies, then sellers are indifferent between applying to a transparent or a non-transparent certifier, since every application to a tier-2 certifier was first rejected by a tier-1 certifier. If the equilibrium features any positive fraction of sellers who go for the safe strategy, then no seller applies to a transparent tier-1 certifier.

Interestingly, there may exist multiple pure equilibria. For example for $d = 0$, the conditions for the safe-strategy and the ambitious-strategy equilibria can be written:

$$\rho S_H(1) \leq \rho S_H(\rho) + (1-\rho)(1-\delta)S_L(\rho) \quad (4)$$

and

$$\rho S_H(1) \geq \rho S_H(0) + (1-\rho)(1-\delta)S_L(0). \quad (5)$$

Indeed, the sellers' certification strategies are *strategic complements*: Ambitious certification strategies depreciate tier-2 certification. A low payoff from being certified by a tier-2 certifier in turn encourages ambitious applications. Focusing on seller welfare W^1 and W^2, Figure 1 depicts the possible equilibrium configurations (the dotted line refers to the outcome under mandated transparency, to be analyzed shortly, and should be ignored for the moment). Eligible beliefs following tier-2 certification range from $\widehat{\rho} = 0$ (all sellers first try a tier-1 certifier, and go to a tier-2 certifier only when rejected) to $\widehat{\rho} = \rho$ (no stigma from tier-2 certification). Under configurations (i) and (ii), one of the strategies (safe and ambitious strategy, respectively) dominates the other, regardless of the stigma attached to tier-2 certification. The equilibrium is then unique. In the third configuration, strategic complementarities lead to three equilibria, two of them stable: all sellers adopt an ambitious strategy ($\widehat{\rho} = 0$ and $W^1(0) \geq W^2(0)$) or all go for the safe strategy ($\widehat{\rho} = \rho$ and $W^2(\rho) \geq W^1(\rho)$). The seller-centric Pareto-dominant equilibrium (yielding the highest welfare among equilibrium outcomes) corresponds to $\widehat{\rho} = \rho$.

FIGURE 1

Proposition 1 *With minimum standard certifiers, certification choices are strategic complements: a seller's choice of an ambitious strategy encourages other sellers to turn to a tier-1 certifier. Furthermore,*

(i) the (Pareto-dominant for the sellers) equilibrium exhibits

- *the ambitious strategy of applying to a non-transparent tier-1 certifier, and then, in case of rejection, to a tier-2 certifier (tiered certification) iff*

$$W^1(\rho) > W^2(\rho), \tag{6}$$

- *the safe strategy of applying directly to a tier-2 certifier otherwise.*

(ii) ambitious strategies are more likely, the lower the probability of fortuitous disclosure (the lower d *is), and the more patient the seller (the higher* δ *is); when* $\delta = 1$ *and* $d = 1$ *ambitious strategies are adopted if and only if the seller is information loving.*

An example of impatient sellers in many American universities is junior faculty members, who are about to come up for tenure.[17] For instance, an assistant professor in the strategy group at a business school may submit a promising empirical analysis to *Management Science*, rather than submitting it to the *American Economic Review*. In part, this choice is driven by the different time frames that the two journals typically have for reviewing papers (on this, see Appendix 3 and the online Appendix). But the junior faculty member may sense that a rejection by a tier-1 certifier would make the track record at the tenure review too thin.[18]

3.2 Regulation of transparency

In reaction to the subprime crisis, the US Treasury chose to require structured investment vehicles to disclose ratings (even unfavorable ones). This section studies whether regulation of disclosure increases welfare in industries in which sellers shop around for certification.[19]

[17] The junior faculty's impatience can reasonably be assumed to be common knowledge, and so we can perform comparative statics with respect to the discount factor (part (ii) of Proposition 1).

[18] If the existing track record makes the tenure decision marginal. Risk-taking by contrast is optimal if the track record is weak.

[19] We focus on governmental regulations. An interesting and related subject of inquiry could be concerned with social regulation (social norms). For example, a social group "regulates" against transparency when ostracizing one of its members who reveals a rejection incurred by another member (in professional or personal matters).

Suppose that a regulator can enforce transparency of applications (this amounts to setting $d = 1$). For given beliefs $\widehat{\rho}$ attached to tier-2 certification, a certification by a tier-2 certifier still yields ("T" refers to "transparency") $W^{2T}(\widehat{\rho}) = W^2(\widehat{\rho})$.

By contrast, application to a tier-1 certifier yields a lower payoff than in the absence of transparency:

$$W^{1T} = W^1(0) < W^1(\widehat{\rho}) \text{ whenever } \widehat{\rho} > 0,$$

The analyis of Section 3.1 applies, with $W^1(\widehat{\rho})$ replaced by $W^1(0)$. Application to a transparent tier-1 certifier (with payoffs as depicted by the dashed horizontal line in Figure 1) is an equilibrium behavior if and only if

$$W^1(0) \geq W^2(\rho).$$

And so if $W^1(0) < W^2(\rho) < W^1(\rho)$, or

$$\rho S_H(1)+\delta(1-\rho)S_L(0) < \rho S_H(\rho)+(1-\rho)S_L(\rho) < \rho S_H(1)+\delta(1-\rho)[(1-d)S_L(\rho)+dS_L(0)],$$

the transparency requirement increases the sellers' welfare: see case (ii) in Figure 1, where $\widehat{\rho} = \rho$ was not an equilibrium under opaque applications and becomes (a Pareto dominant) one under transparency. In the other parameter configurations (cases (i) and (iii) in Figure 1), transparency has no impact on equilibrium outcome and welfare.

Proposition 2 *Transparency weakly improves sellers' welfare.*

Intuitively, transparency makes ambitious strategies less appealing to individual sellers. It thereby eliminates the basic externality associated with ambitious strategies, which depreciate tier-2 certification. Proposition 2 therefore has a natural feel.

Self-Regulation. Relatedly, would tier-1 certifiers agree among each other not to compete on the transparency dimension and to disclose applications? The answer is no, as they would thereby diminish their collective attractiveness.[20]

Buyer welfare. How does transparency impact buyers' welfare? As we have seen, transparency regulation makes a difference only in case (ii) of Figure 1, by killing the ambitious-strategy equilibrium. The issue is thus whether buyers benefit from more or less information. The answer to this question is case-specific. In the first-degree price discrimination illustrations of Examples 1 and 3, buyers have no surplus and so we can confine welfare analysis to that of sellers. In Example 4, either $\rho b_H + (1-\rho)b_L \geq 0$ and then the equilibrium is always a safe-strategy one, or $\rho b_H + (1-\rho)b_L < 0$ and the equilibrium is always the ambitious-strategy one: In either case transparency is irrelevant.

The analysis is more interesting for Example 2 (imperfect price discrimination). In the belief-sensitive-pricing case in Example 2,[21] buyer net surplus in the

[20] To obtain a strict preference for non-transparency, one must assume that certifiers are slightly differentiated (and thus can demand a positive fixed fee).

[21] I.e., when $a_L + b_H > \mu(a_H + b_H)$ and $a_L + b_L < \mu(a_H + b_H)$. The sellers' payoffs in the two potential equilibrium configurations are:

$$W^1 = \rho(a_L + b_H) + \delta(1-\rho)\mu(a_H + b_L)$$

$$W^2 = \begin{cases} a_L + [\rho b_H + (1-\rho)b_L] \text{ for } \rho \geq \rho_0 \\ \mu[a_H + [\rho b_H + (1-\rho)b_L]] \text{ for } \rho < \rho_0. \end{cases}$$

ambitious-strategy and safe-strategy equilibria are:

$$B^1 = \rho\mu(a_H - a_L)$$

$$B^2 = \begin{cases} \mu(a_H - a_L) \text{ for } \rho \geq \rho_0 \\ 0 \text{ for } \rho < \rho_0 \end{cases}$$

respectively. Thus a transparency regulation that moves the equilibrium from ambitious to safe strategies increases (decreases) buyer welfare if $\rho \geq \rho_0$ (if $\rho < \rho_0$). We thus see that while mandated transparency always benefits sellers, it need not benefit buyers. This is a noteworthy observation, in view of the fact that transparency regulation is often heralded as protecting buyers; needless to say, with naive buyers, the case for transparency regulation would be stronger.

4 Extensions

4.1 Transparency and market power in the certification industry

Assume now that the market for tier-i certification is monopolized, while tier-j certifiers are still competitive ($i, j \in \{1, 2\}$).

Market power in tier-1 certification

Suppose first, that the tier-1 monopolist opts for non-transparency. For any fixed fee F charged by the monopoly tier-1 certifier, we can apply the analysis leading to Proposition 1: we only need to replace the function $W^1(\widehat{\rho})$ by the function $W^1(\widehat{\rho}) - F$. We keep assuming that sellers coordinate on the equilibrium that is best for them if there are multiple equilibria, and so from Proposition 1, the monopoly tier-1 certifier attracts the sellers if and only if

$$W^1(\rho) - F > W^2(\rho).$$

Because the monopolist does not want to charge negative fees and thereby lose money, we assume that (6) is satisfied. In equilibrium, sellers opt for the ambitious strategy and apply first to the tier-1 monopolist if and only if $F \leq F^{NT}$ where

$$F^{NT} = W^1(\rho) - W^2(\rho).$$

If (6) is violated, so that we are in case (i) or (iii) of Figure 1, then in equilibrium the sellers Pareto coordinate on the safe strategy for all $F^{NT} \geq 0$ and the tier-1 certfier has no market share. Thus, under non-transparency, the outcome is the same as with a competitive tier-1 industry, except for the monopolist lump-sum payment F^{NT} in case (ii) of Figure 1.

Suppose that instead the monopolist opts for transparency (T). For any fixed fee F charged by the monopolist tier-1 certifier, we can apply the analysis leading to Proposition 1: we only need to replace the function $W^1(\hat{\rho})$ by the function $W^1(0) - F$. When $W^1(0) - W^2(\rho) > 0$—a condition stronger than (6)—then in an equilibrium that is Pareto-undominated for the sellers, sellers opt for an ambitious strategy and apply first to the tier-1 monopolist if and only if $F \leq F^T$ where

$$F^T = W^1(0) - W^2(\rho) < F^{NT}.$$

If $W^1(0) - W^2(\rho) \leq 0$, then the monopolist faces no demand at any non-negative fee $F \geq 0$.

To sum up, a profit maximizing tier-1 certifier strictly prefers non-transparency when (6) is satisfied (case (ii) in Figure 1), and has no market share and is therefore indifferent between transparency or non-transparency when (6) is violated (cases (i) and (iii) in Figure 1).

Proposition 3 *Suppose that tier-2 certification is competitive. A monopoly tier-1 certifier has positive market share if and only if condition (6) is satisfied; the monopolist opts for non-transparency so as to maximize the sellers' incentive to apply for tier-1 certification. Up to a lump-sum transfer, the outcome is exactly the same as for a competitive tier-1 industry.*

Thus, unlike in Lizzeri's (1999) work, disclosure does not hinge on market structure. Note also that this result would hold as well if certifiers did not charge fees and cared only about market share: Regardless of the number of tier-1 certifiers, transparency is a dominated strategy.

Market power in tier-2 certification.

Consider now the opposite case, where the market for tier-1 certification is perfectly competitive, while the market for tier-2 certification is monopolized. Note that in our setup, transparency is a non-issue for a tier-2 certifier. This case is slightly more complicated to analyze, so we simplify the analysis by assuming that $S_H(\hat{\rho}) = S_L(\hat{\rho}) \geq 0$ for all $\hat{\rho}$ and that $d = 0$. Let F be the fixed fee charged by the monopoly tier-2 certifier. We still denote by $W^1(\hat{\rho})$ and $W^2(\hat{\rho})$ the expected gross payoffs when applying to a tier-1 or tier-2 certifier, when certification by a tier-2 certifier delivers reputation $\hat{\rho}$. The only difference is that we are now allowing sellers to refrain from applying to a tier-2 certifier if the price F is so high that their expected surplus is negative. Simmplifying notation, these payoffs can be written as

$$W^1(\hat{\rho}) \equiv \rho S(1) + (1-\rho)\delta \max\{S(\hat{\rho}) - F, 0\}$$

and

$$W^2(\widehat{\rho}) \equiv \max\{S(\widehat{\rho}) - F, 0\}.$$

From Proposition 1, the sellers' Pareto-dominant equilibrium involves an ambitious strategy of applying first to a tier-1 certifier if and only if

$$W^1(\rho) > W^2(\rho).$$

Then the profit of the tier-2 certifier is $(1-\rho)F$ per seller as long as $F \le S(0)$ (in this case, a seller rejected by a tier-1 certifier applies to the tier-2 certifier) and 0 otherwise (in this case, a seller rejected by a tier-1 certifier gives up on certification). On the other hand if

$$W^1(\rho) \le W^2(\rho),$$

which implies $F \le S(\rho)$, then the Pareto-dominant equilibrium involves direct tier-2 applications, and the tier-2 certifier makes a profit of F per seller.

Let $\tilde{F}$ be the fee such that $W^1(\rho) = W^2(\rho)$:

$$\tilde{F} = S(\rho) - \frac{\rho S(1)}{1-(1-\rho)\delta}.$$

Hence we have $W^1(\rho) \le W^2(\rho)$ if and only if $F \le \tilde{F}$.

We assume that the tier-2 certifier does not discount the future and maximizes per-seller profit. There are two cases to consider depending on whether $\tilde{F}$ or $(1-\rho)S(0)$ is greater.

The tier-2 certifier sets $F = S(0)$ if $(1-\rho)S(0) > \tilde{F}$ and $F = \tilde{F}$ otherwise. In the first case, the Pareto-dominant equilibrium involves an ambitious strategy, and in the second case, a safe strategy. Turning to the competitive case (which is equivalent to $F = 0$), we see that the Pareto-dominant equilibrium involves an ambitious strategy if $\tilde{F} < 0$ and a safe strategy otherwise. Hence a monopolized tier-2 certification market leads to an ambitious strategy equilibrium for a larger set of parameter values compared to the perfectly competitive case. This result is to be contrasted with Proposition 1. In essence, tier-2 market power makes tier-2 certification less attractive and favors ambitious strategies.

Proposition 4 *Suppose that tier-1 certification is competitive, that* $S_H(\cdot) = S_L(\cdot)$ *and that* $d = 0$*. With a monopoly tier-2 certifier, the outcome is an ambitious strategy equilibrium for a larger set of parameter values than with competitive tier-2 certifiers.*

The key feature leading to lack of transparency under perfect competition in the certification industry is that the sellers choose their certifiers, not that they pay them.

4.2 Seller heterogeneity

Next we allow for seller heterogeneity when tier-1 certification is competitive or monopolized. Seller heterogeneity raises the possibility of sorting through contracts that are differentiated in the payment terms and in the ensuing reputation pattern. We assume that ρ lies in the continuum $[0, 1]$ and that only the seller knows ρ. In equilibrium, sellers with $\rho > \rho^*$, for some cutoff ρ^*, go for tier-1 certification and sellers with $\rho < \rho^*$ for the safe strategy. As we will see, non-transparency is optimal for tier-1 certifiers for three reasons. The first two reasons are the same as in the case of seller homogeneity: transparency reduces the individual seller's utility as well as the stigma from tier-2 certification. The third reason, specific to heterogeneity, also argues against transparency: among types selecting tier-1 certification, transparency hurts high types less than low types and so is a bad screening device to extract the high types' rent.

For notational simplicity, we assume that $d = 0$.

Competitive certification industry. In equilibrium, types $\rho > \rho^*$ go to a tier-1 certifier while types $\rho < \rho^*$ do not, where the cutoff ρ^* as well as the probability $\widehat{\rho}$ of a high type given tier-2 certification are given by:

$$\rho^* S_H(1) + \delta(1-\rho^*) S_L(\widehat{\rho}) = \rho^* S_H(\widehat{\rho}) + (1-\rho^*)\, S_L\,(\widehat{\rho})$$

and

$$\widehat{\rho} = \frac{\int_0^{\rho^*} \rho g(\rho) d\rho}{G(\rho^*) + \int_{\rho^*}^1 (1-\rho) g(\rho) d\rho}.$$

These two equations yield two monotonically increasing mappings between ρ^* and $\widehat{\rho}$, confirming the existence of strategic complementarities. There can be multiple equilibria: equilibria with high values of $\hat{\rho}$ and ρ^* and equilibria with low values of $\hat{\rho}$ and ρ^*. The former correspond to the safe strategy equilibria, and the latter to the ambitious strategy equilibria of the homogenous seller case of Section 3.

As in Section 3, the analysis of equilibria is not affected by entry of transparent tier-1 certifiers. These certifiers do not attract any sellers in equilibrium.

Monopoly tier-1 certifier. A monopoly tier-1 certifier builds an incentive-compatible mechanism $\{F(\rho), \widehat{\rho}(\rho)\}$, where $F(\rho)$ is the fee demanded from type ρ and $\widehat{\rho}(\rho)$ is the ex-post reputation in case of rejection: $\widehat{\rho}(\rho) = 0$ in case of disclosure/transparency and $\widehat{\rho}(\rho) = \widehat{\rho}$ (the general reputation in the absence of disclosure) otherwise.

Type ρ's utility is

$$W(\rho) = \rho S_H(1) + \delta(1-\rho) S_L(\widehat{\rho}(\rho)) - F(\rho)$$

and so incentive compatibility can be written as

$$\frac{dW(\rho)}{d\rho} = S_H(1) - \delta S_L(\widehat{\rho}(\rho)) > 0.$$

There is a cutoff ρ^* such that sellers apply to the tier-1 certifier if and only if $\rho > \rho^*$, where $W(\rho^*) = \rho^* S_H(\widehat{\rho}) + (1-\rho^*)\, S_L(\hat{\rho})$. Then $\hat{\rho}$ is determined by

$$\widehat{\rho} = \frac{\int_0^{\rho^*} \rho g(\rho) d\rho}{G(\rho^*) + \int_{\rho^*}^1 (1-\rho) g(\rho) 1\{\hat{\rho}(\rho) = \hat{\rho}\} d\rho}. \tag{7}$$

where $1\{\hat{\rho}(\rho) = \hat{\rho}\}$ is the indicator function taking the value 1 if $\hat{\rho}(\rho) = \hat{\rho}$ and 0 otherwise.

Given ρ^*, we integrate the incentive compatibility constraint to find for $\rho > \rho^*$

$$W(\rho) = \rho^* S_H(\widehat{\rho}) + (1-\rho^*)\, S_L(\hat{\rho}) + \int_{\rho^*}^{\rho} [S_H(1) - \delta S_L(\widehat{\rho}(\tilde{\rho}))]\, d\tilde{\rho}. \tag{8}$$

We can then write down the certifier's profit as follows

$$\int_{\rho^*}^1 [\rho S_H(1) + \delta(1-\rho) S_L(\widehat{\rho}(\rho)) - W(\rho)]\, g(\rho) d\rho,$$

which we can rewrite using (8) as

$$\int_{\rho^*}^1 [\rho^* S_H(1) - \rho^* S_H(\widehat{\rho}) - (1-\rho^*)\, S_L(\hat{\rho}) + \delta(1-\rho) S_L(\widehat{\rho}(\rho))]\, g(\rho) d\rho \tag{9}$$
$$+ \int_{\rho^*}^1 \int_{\rho^*}^{\rho} \delta S_L(\widehat{\rho}(\tilde{\rho})) d\tilde{\rho} g(\rho) d\rho.$$

The certifier's problem is then to maximize (9) over ρ^* and $\{\hat{\rho}(\rho)\}_{\rho \geq \rho^*}$, subject to (7). A simple inspection of this maximization shows that it is optimal to set $\widehat{\rho}(\rho) = \widehat{\rho}$ (no transparency) for all $\rho \geq \rho^*$. The fee $F(\rho) = \rho S_H(1) + \delta(1-\rho) S_L(\widehat{\rho}) -$

$W(\rho)$ is independent of ρ (non-responsiveness).

Proposition 5 *Non-transparency obtains under heterogeneous sellers whether tier-1 certification is competitive or monopolized. Furthermore, with heterogeneous sellers, a monopoly tier-1 certifier charges a type-independent fee.*

It is optimal for a tier-1 monopolist to adopt non-transparency for three reasons. First, this policy maximizes the stigma associated with tier-2 certification; second, among tier-1 customers, transparency does not help screen high types since it hurts them less than lower types; finally, transparency directly reduces the utility of tier-1 customers by lowering their utility in case of rejection to $S_L(0)$.

4.3 Multi-tier certification

Assume now that certifiers can, at cost $c \geq 0$, provide a fine grade if they choose so (which, for a competitive certification industry, is equivalent to the sellers wanting a fine grade). We maintain the assumption that $d = 0$ for expositional simplicity. In the same way they do not want to disclose unsuccessful applications, tier-1 certifiers do not gain by transforming themselves into multi-tier certifiers. The question is then whether tier-2 certifiers are displaced by multi-tier certifiers and how this affects the sellers' incentive to apply to tier-1 certifiers.

The broad intuition, which we develop in more detail below, goes as follows: Sellers who would otherwise have applied directly to a tier-2 certifier can avoid the adverse-selection stigma by turning to a multi-tier certifier. This stigma avoidance involves cost c and comes at a further cost if sellers are information averse. If sellers are information loving or neutral, and the cost of fine grading is small, multi-tier certification drives out tier-2 certifiers; it also drives out tier-1 certifiers as resubmission after a rejection by a tier-1 certifier involves a delay and cannot prevent the buyers from knowing that quality is not high. Thus, if fine grading is costless, minimum-standard certification can survive only if sellers are information averse.

Consider first an ambitious-submission equilibrium ($x = 1$) under minimum-standard certification (Section 3). Sellers obtain $\rho S_H(1) + \delta(1-\rho)S_L(0)$. But they can avoid discounting and obtain $\rho S_H(1) + (1-\rho)S_L(0) - c$ by turning to a multi-tier certifier directly. The tiered-certification equilibrium therefore requires, besides condition (6), that

$$\rho S_H(1) + \delta(1-\rho)S_L(0) \geq \rho S_H(1) + (1-\rho)S_L(0) - c$$

$$\iff c \geq \underline{c} \equiv (1-\delta)(1-\rho)S_L(0).$$

If $c < \underline{c}$, tier-2 certifiers overcome the stigma they face in an ambitious-strategy equilibrium by converting to multi-tier certification.

Second, consider a safe-strategy equilibrium ($x = 0$), and so condition (1) obtains. This equilibrium is robust to the introduction of full-grading if and only if furthermore

$$\rho S_H(\rho) + (1-\rho)S_L(\rho) \geq \rho S_H(1) + (1-\rho)S_L(0) - c$$

$$\iff c \geq \overline{c} \equiv \overline{c} \equiv \rho[S_H(1) - S_H(\rho)] - (1-\rho)[S_L(\rho) - S_L(0)].$$

Note that when $c = 0$, this condition holds if and only if the sellers are information averse.[22]

[22] For the sake of completeness, we can consider a mixed equilibrium ($0 < x < 1$). A necessary and sufficient condition for this equilibrium to be robust to the introduction of fine grading is that the sellers who apply directly to a tier-2 certifier do not find it advantageous to go for a full grade:

$$\rho S_H(\widehat{\rho}(x)) + (1-\rho)S_L(\widehat{\rho}(x)) \geq \rho S_H(1) + (1-\rho)S_L(0) - c.$$

To sum up, sellers resort to multi-tier grading when its cost c is low, when sellers are impatient (δ is low), and when sellers are information neutral or loving.

Proposition 6 *Multi-tier grading occurs for a larger set of parameters, the lower its cost* c*, the more impatient the sellers (the lower* δ*), and the less information-averse the sellers are (the higher* $\rho[S_H(1) - S_H(\rho)] - (1-\rho)[S_L(\rho) - S_L(0)]$*).*

Entry attempts into multi-tier grading seem to resonate with academic experience. Illustrations include fine grading by bepress and the proliferation of prizes offered by tier-2 journals (and not by tier-1 journals). Alternatively, multi-tier grading can piggyback on an existing tier-1 certifier. For example, the American Economic Association launched the *American Economic Journals*; to be certain, the AEA is careful about not describing the AEJs as lower tier relative to the *American Economic Review*, and the process is a bit different from a multi-tier process, but the fundamentals driving the introduction of the AEJs are the same as those developed in this section. *Nature* operates a very similar system via its set of "Nature Research Journals," of which *Nature Biotechnology* is probably one of the best known.[23]

Proposition 6 may also shed some light on rating agencies' practice of fine grading. As we observed in Example 3 (Section 2), bond ratings not only certify the quality of an issue but also allow matching between securities and buyers.

[23] Both for the *American Economic Review* and for *Nature*, the process is similar. For example, If a paper submitted to *Nature* is close to being publishable (but not quite so), and a good fit with a *Nature* research journal, then the Nature editors add a paragraph in their decision letter encouraging the author to send it to the appropriate *Nature* research journal. If the author does choose to submit to a research journal after such a referral (the research journal editors are cc'ed on letters containing referrals), the *Nature* editor will be notified, and will send the editor of the research journal the Nature referee reports. The *Nature* editor will also sometimes talk to the editor of the research journal. The research journal editor (a) may accept the paper without further review, (b) may ask for changes, followed by either an acceptance or re-refereeing, or (c) may just reject the paper.

This matching dimension became more important in the mid 1970s, when broker-dealers' regulatory assessment of solvency (and then insurers', pension funds', and, with Basel II, banks') started to make use of ratings, creating a strong demand for high-quality liquid claims. The matching dimension is captured here by sellers becoming more information loving. The mid-1970s coincidentally were a turning point in the business model of rating agencies, which switched to the issuer-pays mode.

Remarks: Our assumption that certifiers can commit to a given disclosure policy may be a bit stretched in the case of multi-tier grading. Suppose that such a commitment is enforced by reputational concerns, and consider a tier-2 certifier trying to break a tiered-certification equilibrium by converting into a multi-tier grade certifier. If sellers do not believe in this strategy, the certifier is deprived of high types and cannot (and has no incentive to) develop a reputation for accurate, fine grading. As we earlier indicated, we leave foundations of commitment for future research, but we note that our commitment assumption may be more problematic for some forms of certification than for others.[24]

Proposition 6 focuses on a competitive certifying industry. In an online appendix, we consider a monopoly certifier who can costlessly engage in fine grading; we perform a mechanism design exercise and shows how efficient disclosure relates to the sellers' information aversion.

[24] We can, however, capture this idea through the following reduced form: Suppose that each certifier secretly chooses between spending 0 and spending c per review (say, by recruiting talented employees), and announces publicly its certification strategy (tier-1, tier-2, multi-tier); and that it incurs a finite penalty for incorrect rankings. No certifier has an incentive to invest in the cost c per review if sellers choose an ambitious strategy and believe that certifiers do not invest in the extra cost.

4.4 Further extensions

The appendix analyzes three other extensions. Appendix 1 endogenizes the quality of the seller's product. Quality depends on the seller's costly effort e. Higher effort increases the probability that the quality is high, but does not change the probability that the quality is abysmal. We completely characterize the equilibrium. We show that quality is always higher in an ambitious-strategy equilibrium than in a safe strategy equilibrium. Moreover, we show that transparency weakly improves sellers' welfare, and weakly reduces quality. When it does so strictly, a safe-strategy equilibrium with low quality investment replaces an ambitious-strategy equilibrium with high quality investment.

Appendix 2 maintains the assumption that sellers select certifiers, but introduces the possibility that certifiers charge buyers (this requires that the certifier be able to prevent buyers from reselling the information to each other). It shows that the analysis of this paper is robust to charges levied by certifiers on buyers.

Appendix 3 allows for a class of quick turn-around tier-1 certifiers. Shorter certification lags result in reduced accuracy in the form of type-I and type-II errors.[25] We completely characterize the equilibrium. We derive explicit conditions under which the equilibrium features sellers first applying to a quick turn-around tier-1 certifier and then to a tier-2 certifier in case of rejection. Interestingly, the structure of equilibria depends on whether the decisions to apply to a quick turn-around tier-1 certifier instead of a tier-1 certifier are strategic complements or substitutes. We show that they are strategic complements when the probability of a type-1 error is high, the probability of a type-2 error is small, and the turn-around advantage of the quick turn-around certifier is high. In an online appendix, we further show

[25] Of course quick turn-around need not be associated with lower quality and may just result from superior effort or norms. But choices become meaningful when within a given category of certifiers, turn-around-time and accuracy covary negatively.

that the more competitive the industry, the more likely it is that certifiers offer a quick (slower) turn-around time if certifiers maximize market share (profit).

5 Summary and conclusion

Certifiers such as journals, rating agencies, standard setting bodies and providers of standardized tests play an increasingly important role in our market economies. This paper makes an initial attempt at understanding how the certification industry caters to the sellers' demand through strategies such as the non-disclosure of rejections, and analyzes the welfare implications of such policies.

In the absence of regulation, certifiers do not publicize rejected applications. First, such disclosure reduces the individual seller's (the certifier's clients) utility. Second, when there is market power in tier-1 certification, the certifier is eager to increase the stigma from tier-2 certification, and thereby soften competition, by not disclosing rejections. Third, and again under tier-1 market power, but when sellers are heterogeneous, disclosure does not help capture the rent of the most confident sellers, as the latter are less affected by disclosure than less-confident ones.

On the positive side, we also examined when sellers are willing to take the risk of applying to a tier-1 certifier. This willingness hinges on the behavior of other sellers (which affects the stigma associated with a tier-2 acceptance, leading to strategic complementarities), the discount factor (which impacts the cost of an ambitious submission strategy), and sellers' information aversion (which determines the reputation-risk tolerance).

On the normative side, sellers' gaming of the certification process involves costs: delay (or, in a variant of our model, duplication of certification costs) and possibly excessive information exposure; as self-regulation does not promote transparency, these costs were shown to provide a role for public regulation. We showed that transparency regulation always benefits sellers, but need not benefit buyers.

Finally, we investigated a number of extensions/robustness checks by allowing multi-tier certification, fees charged to buyers, endogenous choice of quality by sellers, and endogenous choice of turn around by certifiers.

Turning back to Table 1, it is not surprising in light of our theoretical predictions that the bulk of the entries are under the opaque heading. State licensing examinations may be fundamentally different due to the presence of regulatory dicta; accordingly, "sellers" cannot choose their certifier. Entry-level examinations exhibit transparency, but may or may not exhibit fine grading. These features may reflect the power imbalance between the buyers (say, colleges) and sellers (would-be students). In this instance, it is the buyers rather than the sellers who choose certifiers, which probably explains the unusual entry in Table 1.[26] Finally, and also consistent with our theory, it is not surprising that in situations where we would anticipate that information aversion would be greatest (e.g., an undergraduate or MBA student going on the job market, an entrepreneurial firm going public), we see minimum standard certification rather than a fine-grained scheme.

[26] Top schools want to be matched with top students. They therefore have an incentive to demand transparency as well as tier-1 certification or a fine grading.

References

Biglaiser, G., “Middlemen as Experts,” *Rand Journal of Economics*, 24 (1993), pp. 212–223.

Bolton, P., Freixas, X., and Shapiro, J. “The Credit Rating Game,” *Journal of Finance*, 67 (2011), pp. 85–112.

Bolton, P., Santos, T., and Scheinkman, J. A. "Cream Skimming in Financial Markets," mimeo, Columbia University and Princeton University, 2011.

Bouvard, M., and Levy, R. “Humouring Both Parties: A Model of Two-Sided Reputation,” mimeo, Toulouse School of Economics, 2009.

Brandt, R., and Weisel, T. *Capital Instincts: Life As an Entrepreneur, Financier, and Athlete*. New York, John Wiley, 2003.

Coffee, J.C., Jr. “Turmoil in the U.S. Credit Markets: The Role of the Credit Rating Agencies,” Testimony Before the U.S. Senate Committee on Banking, Housing and Urban Affairs, April 22, 2008.

Faure-Grimaud, A., Peyrache, E., and Quesada, L. “The Ownership of Ratings,” *Rand Journal of Economics*, 40 (2009), pp. 234–257.

Gill, D., and Sgroi, D. “The Superiority of Tough Reviewers in a Model of Simultaneous Sales,” DAE Working paper 335, Cambridge University, 2004.

Jeon, D.S., and Rochet, J.C. “The Pricing of Academic Journals: A Two-Sided Market Perspective,” *American Economic Journal: Microeconomics*, 2 (2010), pp. 222–255.

Lerner, J., and Tirole, J. “A Model of Forum Shopping,” *American Economic Review*, 96 (2006), pp. 1091–1113.

Lizzeri, A. "Information Revelation and Certification Intermediaries," *Rand Journal of Economics*, 30 (1999), pp. 214–231.

Mathis, J., McAndrews, J. and Rochet, J.C. "Rating the Raters: Are Reputation Concerns Powerful Enough to Discipline Rating Agencies?" *Journal of Monetary Economics*, 56 (2009), pp. 657–674.

McCabe, M.J., and Snyder, C. "Open-Access and Academic Journal Quality," *American Economic Review Papers and Proceedings*, 95 (2005), pp. 453–458.

———- "Academic Journal Prices in a Digital Age: A Two-Sided Market Approach," *B.E. Journal of Economic Analysis and Policy* (Contributions). January, vol. 7 (2007), article 2.

———- "The Economics of Open-Access Journals," mimeo, University of Michigan and Dartmouth College, 2010.

Morrison, A., and White, L. "Crises and Capital Requirements in Banking," *American Economic Review,* 95 (2005), pp. 1548–1572.

Partnoy, F. "How and Why Credit Rating Agencies are Not Like Other Gatekeepers," in Yasuyuki Fuchita and Robert E. Litan, eds., *Financial Gatekeepers: Can They Protect Investors*?, Washington: Brookings Institution Press and the Nomura Institute of Capital Markets Research, 2006.

Rimer, S. "SAT Changes Policy, Opening Rift With Colleges," *New York Times*, December 30, 2008, p. A14.

Sangiorgi, F., Sokobin, J., and Spatt, C. "Credit-Rating Shopping, Selection and the Equilibrium Structure of Ratings," mimeo, Stockholm School of Economics, and NBER, 2009.

Skreta, V., and Veldkamp, L. "Ratings Shopping and Asset Complexity: A Theory of Ratings Inflation," *Journal of Monetary Economics*, 56 (2009), pp. 678–695.

U.S. Securities and Exchange Commission. *Report on the Role and Function of Credit Rating Agencies in the Operation of the Securities Markets, As Required by Section 702(b) of the Sarbanes-Oxley Act of 2002*, Washington: Government Printing Office, and associated testimony and exhibits (particularly that of Sean J. Egan), 2003.

U.S. Securities and Exchange Report, Office of Compliance Inspections and Examinations and Office of Economic Analysis. *Summary Report of Issues Identified in the Commission Staff's Examinations of Select Rating Agencies*, Washington, SEC, 2008.

Appendix 1 (endogenous quality)

This Appendix shows that our analysis is unchanged when the choice of quality depends on the equilibrium of the certification process. Suppose that quality depends on the seller's investment effort $e \in [\underline{e}, \overline{e}]$. We are interested in modeling a dimension of effort that affects the likelihood of a high quality outcome but does not change the probability of an abysmal outcome. It is reasonable to think that those margins respond to different forms of investment, and that for some of the examples that we have in mind, the latter margin would be quite inelastic.[27] Hence our focus on the former.

Let q be the probability that a product is not abysmal. A higher effort increases the probability of the high quality $\rho(e)$ outcome conditional on a non-abysmal outcome. Let $\psi(e)$ denote the disutility of effort. We assume that $\rho(e)$ is increasing and concave in e and that $\psi(e)$ is increasing and convex in e with $\rho'(\underline{e}) = +\infty$ and $\psi'(\underline{e}) = 0$. To simplify the analysis, we also assume that $S_L(.) = S_H(.)$ (as in Examples 1 through 3), and that $d = 0$.

We define two ex-ante payoff functions $\mathcal{W}^1$ and $\mathcal{W}^2$ as follows:

$$\mathcal{W}^1(\hat{\rho}) \equiv \max_e \{q[\rho(e) S(1) + \delta(1-\rho(e)) S(\hat{\rho})] - \psi(e)\}$$

and

$$\mathcal{W}^2(\hat{\rho}) \equiv \max_e \{qS(\hat{\rho}) - \psi(e)\}.$$

Let $e^1(\hat{\rho})$ and $e^2(\hat{\rho})$ be the solutions of the maximization problems underlying $\mathcal{W}^1$ and $\mathcal{W}^2$. Clearly, $e^2(\hat{\rho}) = \underline{e}$.

Lemma 7 *We have* $\frac{d\mathcal{W}^2(\hat{\rho})}{d\hat{\rho}} > \frac{d\mathcal{W}^1(\hat{\rho})}{d\hat{\rho}}$ *for all* $\hat{\rho}$.

[27] More generally, the analysis extends straightforwardly to a small elasticity of abysmal quality to effort.

Proof. By the envelope theorem,

$$\begin{aligned}\frac{d\mathcal{W}^1(\hat{\rho})}{d\hat{\rho}} &= q\delta\left(1-\rho\left(e^1(\hat{\rho})\right)\right)\frac{d\mathcal{S}(\hat{\rho})}{d\hat{\rho}}\\ \frac{d\mathcal{W}^2(\hat{\rho})}{d\hat{\rho}} &= q\frac{d\mathcal{S}(\hat{\rho})}{d\hat{\rho}}\end{aligned}$$

The result follows immediately. ■

There are two potential equilibria. The ambitious strategy equilibrium effort level e^{1*} and the safe-strategy equilibrium effort level e^{2*} are determined by the following equations:

$$e^{1*} = e^1(0) > \underline{e} = e^{2*}.$$

The safe strategy is an equilibrium if and only if

$$\mathcal{W}^2(\rho(\underline{e})) \geq \mathcal{W}^1(\rho(\underline{e}))$$

while the ambitious strategy equilibrium is an equilibrium if and only if

$$\mathcal{W}^1(0) \geq \mathcal{W}^2(0).$$

From Lemma 1, an equilibrium always exists. The safe and risky strategy equilibria coexist over a range of parameters. When there are multiple equilibria, we select the ex-ante Pareto dominant equilibrium. Let $\mathcal{W}^{1T} \equiv \mathcal{W}^1(0)$ denote the sellers' welfare under mandated transparency. The analysis is then identical to the case where effort is exogenous, with ρ replaced by $\rho(\underline{e})$ and W^1, W^{1T}, and W^2 replaced by $\mathcal{W}^1$, $\mathcal{W}^{1T}$, and $\mathcal{W}^2$. In particular, transparency weakly improves sellers' welfare. When it does so strictly, a safe-strategy equilibrium with low-quality investment replaces an ambitious-strategy equilibrium with high-quality investment.

Proposition 8 *When quality is endogenous,*

(i) quality is always higher in an ambitious-strategy equilibrium than in a safe-strategy one;

(ii) transparency weakly improves sellers' welfare and weakly reduces quality.

Appendix 2 (buyer-pay paradigm)

We could alternatively assume that certifiers can monitor that buyers do not communicate the ratings among each other, and so that they can charge buyers rather than sellers (how much they can charge depends on the context, as the buyers' willingness to pay depends on the anticipated pricing strategy of the sellers). Imagine now that certifiers are perfectly competitive profit maximizers and can charge buyers but not sellers. Suppose for simplicity that when deciding whether to buy the rating, buyers do not yet know their type. Then the payment to the certifier is a lump-sum payment, equal to the buyers' expected net surplus, and has no influence on seller payoffs $S_i(\widehat{\rho})$. Consider, for example, Example 2 of Section 2.2 (the logic extends more generally). When condition (6) is satisfied, and letting as earlier $B(\widehat{\rho})$ denote buyer surplus, $B(\widehat{\rho}) = \mu(a_H - a_L)$ if $\widehat{\rho} \geq \rho_0$, $B(\widehat{\rho}) = 0$ if $\widehat{\rho} < \rho_0$. Tier-1 certifiers can charge fee $F_B = B(1)$ for the disclosure of a successful tier-1 application. Tier-2 certifiers can charge a fee $F_B = B(0) = 0$ for the disclosure of a successful tier-2 application. In this ambitious-strategy equilibrium, buyers know that a tier-2 certification means that the seller was previously rejected by a tier-1 certifier, and thus neither tier-1 nor tier-2 certifiers are able to charge anything for the disclosure of a failed tier-1 application. In this case, disclosing this information for free to buyers is indifferent to sellers. Suppose instead that condition (6) is violated. Then tier-2 certifiers can charge $F_B = B(\rho)$ for the disclosure of a successful tier-2 application. There is then no way for a tier-1 certifier to make a profit and attract sellers, whether or not they disclose failed applications and charge for it or not. The analysis of the equilibrium is therefore completely identical to the case where certifiers charge sellers and not buyers. The only difference is that certifiers appropriate the buyer's surplus that is not appropriated by sellers.

For example, academic journals have traditionally charged the buying side. They bundled, however, the certification and distribution function. The distribution function nowadays can be performed through web sites and web repositories (although journals try to keep the two activities bundled through requirements not to keep papers posted once they are accepted). The recent advocacy in favor of open access publishing (charging authors through submission and/or publication fees, rather than readers) may accelerate this unbundling. An interesting literature (e.g., McCabe-Snyder 2005, 2007,2010 and Jeon-Rochet 2010) analyzes certification from the point of view of two-sided markets theory. In particular, it looks at when academic journals should charge readers or authors, and how the quality of certification is affected by this choice. By way of contrast, the issues of transparency and sequential certification remain to be investigated in full generality in this context.

To sum up, because the analysis rests entirely on the seller's surplus $S_i(\widehat{\rho})$, nothing is altered by introducing a buyer fee,[28] in the same way the analysis was shown to carry over to the case of a monopolistic certifier charging a seller fee. The non-transparency result is thus very robust.

[28] The certifier then obtains a rent. This rent can be dissipated either through free entry when there is a fixed entry cost into the certification industry (monopolistic competition) or through a subsidy to sellers for an exclusive certification (as emphasized by two-sided markets theory).

Appendix 3 (quick turn-around)

We now assume that tier-1 certifiers choose their certification delays so as to attract sellers. Shorter lags may increase the certification cost (here normalized at 0) or (as modeled here) result in reduced accuracy. We assume that quick turn-around certifiers make type-I and type-II errors. They receive a high signal H with probability $1 - z_H$ if the actual quality is high, and z_L if it is low, where $1 - z_H > z_L$. Thus, they act as tier-1 certifiers with noisy signals. To capture the idea that short turn-around times benefit the sellers, we assume that a quick turn-around certification takes less time (and therefore is subject to discount factor $\widehat{\delta} > \delta$), while both tier-1 and tier-2 certification take a full period.[29] To ensure that these quick turn-around certifiers are not able to supersede the tier-2 certifiers in ensuring that quality is not abysmal, we assume that they look for an H type (only H and L types can get the H signal) but in the absence of an H signal, cannot rule out the abysmal quality.

A seller who is rejected by a quick turn-around certifier can apply to a tier-2 certifier without losing as much time as if he had been rejected by a tier-1 certifier. We will make assumptions so that it is never optimal to turn directly to a tier-2 certifier, and that it is never optimal to turn to a quick turn-around certifier after a rejection either by a tier-1 certifier or by a quick turn-around certifier. We further assume that $d = 0$, and that $S_H(\widehat{\rho}) = S_L(\widehat{\rho}) \equiv S(\widehat{\rho})$ for all $\widehat{\rho}$, so as to simplify the analysis.

[29] In order to avoid integer problems (and the concomitant possibility that the date of product introduction reveal the strategy), one must assume in this section that sellers arrive in continuous time (but the certification length is still discrete).

Alternatively, we could assume that quick turn-around results in a random certification length of either 0 or 1, while regular certification always takes 1 period.

Let

$$\rho^{+} \equiv \frac{\rho(1 - z_H)}{\rho(1 - z_H) + (1 - \rho)z_L}$$

denote the posterior belief following an H signal by a quick turn-around certifier. Such a signal is good news for the quality of the product, i.e. $\rho^{+} > \rho$, since $1 > z_H + z_L$.

Our first assumption adapts condition (6) to the setup of this section. It ensures that no safe-strategy equilibrium exists, as sellers would then rather choose to try a tier-1 certifier first:

$$\rho S(1) + (1 - \rho)\,\delta S(\rho) > S(\rho). \tag{10}$$

Our second assumption is sufficient to ensure that after a rejection by a tier-1 certifier, a seller prefers to apply to a tier-2 certifier than to try a quick turn-around certifier (and then a tier-2 certifier in case of rejection by the quick turn-around certifier):

$$\delta S(0) \geq \hat{\delta}\left(z_L S(1) + (1 - z_L)\,\delta S(0)\right). \tag{11}$$

Our third assumption is sufficient to ensure that after a rejection by a quick turn-around certifier, a seller prefers to apply to a tier-2 certifier than to a tier-1 certifier (and then a tier-2 certifier in case of rejection by the tier-1 certifier):

$$\delta S(0) \geq \delta[\widehat{\rho}_2(1)S(1) + \delta[1 - \widehat{\rho}_2(1)]S(0)], \tag{12}$$

where $\widehat{\rho}_2(1)$, the posterior beliefs following a rejection by a quick-turn-around certifier, is defined below. Note that the left-hand-side of this equation corresponds to the most pessimistic beliefs possible—this is why is condition is sufficient but not necessary.

Last, we ensure that a seller does not want to turn to another quick turn-around certifier after being rejected by one. A sufficient condition for the absence of such repeated attempts is that false positives be perfectly correlated among quick turn-around certifiers, and so a failed attempt to be certified by such a certifier does not lead to other attempts.

Given these assumptions, the only relevant strategic consideration is whether to apply to a quick turn-around certifier or to a tier-1 certifier. Denote by y the fraction of applicants who opt for a quick turn-around certification rather than tier-1 certifiers.

Let $\widehat{\rho}_2 = \widehat{\rho}_2(y)$ denote the posterior beliefs following tier-2 certification:

$$\widehat{\rho}_2\,(y) = \frac{y\rho z_H}{y\rho z_H + y(1-\rho)(1-z_L) + (1-y)(1-\rho)}.$$

We necessarily have $\rho^+ > \rho > \widehat{\rho}_2\,(y)$. Both false positives and false negatives improve the pool of applications to tier-2 certifiers and decrease the stigma associated with tier-2 certification. As long as $z_H > 0$, $\widehat{\rho}_2\,(y)$ increases with y as the stigma associated with tier 2 certification decreases.

For a given probability y, sellers turn to a certifier with low turn-around time rather than to a tier-1 certifier if and only if $\Psi(y) \geq 0$ where:

$$\begin{aligned}\Psi(y) \quad &= \quad \hat{\delta}[\rho(1-z_H) + (1-\rho)z_L]S(\rho^+) + [\rho z_H + (1-\rho)(1-z_L)]\hat{\delta}\delta S(\widehat{\rho}_2(y)) \\ &\quad\quad -\delta[\rho S(1) + \delta(1-\rho)S(\widehat{\rho}_2(y))].\end{aligned}$$

The sign of $\Psi'(y)$ determines whether the choices between tier-1 certification and quick turn-around certification are strategic complements (positive sign) or substitutes (negative sign). Decisions are strategic complements if and only if

$$\rho z_H + (1-\rho)(1-z_L) \geq \frac{\delta}{\hat{\delta}}(1-\rho). \tag{13}$$

The left-hand side of (13) is the probability of being rejected when applying to a quick turn-around certifier. The right-hand side of (13) is the discounted probability of being rejected by a tier-1 certifier. Increasing y reduces the stigma of applying to a tier-2 certifier, which impacts the payoff of both the tier-1 certification strategy and the quick turn-around application strategy in proportion to these probabilities. The higher z_H, the lower z_L and the lower δ, the more likely is (13) to be verified.

When (13) holds, then there can be multiple equilibria. This occurs when the following additional conditions hold:

$$\Psi(0) < 0 < \Psi(1). \tag{14}$$

If there are multiple equilibria, the equilibrium where all sellers first turn to quick turn-around certifiers has higher seller welfare. Indeed, combining a revealed preference argument ($\Psi(1) > 0$) and the fact that $\rho S(1) + \delta(1-\rho)S(\hat{\rho}_2(1)) > \rho S(1) + \delta(1-\rho)S(0)$ yields the result. We maintain the maximization of seller welfare as our selection criterion, and so the economy will find itself in the quick turn-around equilibrium as long as $\Psi(1) > 0$. Again, if multiple equilibria coexist, sellers are better off in the one with the least stigma attached to tier-2 certification.

When (13) is violated, the equilibrium is unique, and may entail mixed strategies. If $\Psi(1) \geq 0$ (and hence $\Psi(0) > 0$), then the equilibrium involves quick turn-around certification. When $\Psi(0) \leq 0$ (and hence $\Psi(1) < 0$), then the equilibrium involves tier-1 certification. When $\Psi(1) < 0 < \Psi(0)$, then the equilibrium involves mixed strategies.

Proposition 9 *Suppose that* $0 < z_H < 1 - z_L$ *and that (10) through (12) hold. If (13) holds, then the equilibrium involves quick turn-around certification if* $\Psi(1) \geq 0$ *and tier-1 certification otherwise. If (13) is violated, then the equilibrium involves quick turn-around certification when* $\Psi(1) \geq 0$*, tier-1 certification when* $\Psi(0) \leq 0$*, and mixed strategies otherwise.*

Market structure and quick turn-around

The online appendix analyzes how market structure affects the emergence of quick turn-around certification versus tiered certification. More specifically, it maintains the assumption that the market for tier-2 certifiers is perfectly competitive, and analyzes the impact of the degree of competition among tier-1 certifiers.

We show that the effect of competition depends on its nature, namely whether certifiers choose prices to maximize profit or else do not set prices and maximize market shares. We find conditions under which competition enhances quick turn-around certification when certifiers compete in market shares and not in prices.