

Optimal Quality Certification

By ANDRIY ZAPECHELNYUK*

Quality certification not only informs consumers, but also stimulates producers to supply better quality products. We study a problem of quality certification in a moral hazard setting. We show that, under standard assumptions, simple certification systems, such as quality assurance rule and pass-fail rule, are optimal. Our solution method involves interpreting the certification problem as a delegation problem.

JEL: D82, D86, D45

Keywords: Certification, Bayesian persuasion, information disclosure, information design, delegation, moral hazard, career concerns

Consumers often do not observe the true quality of products, thus relying on indirect methods of ascertaining their worth. One of such methods is *certification*, where an independent certifier verifies conformity of a product to specified standards. Certification is often mandatory and regulated by the government, especially when it concerns health, safety, and wellbeing. For example, car manufacturers must disclose conformity to crash safety standards. Childcare and educational organizations obtain accreditation from an office on standards in education. Restaurants and catering businesses obtain food hygiene certificates from a food standards agency. Food manufacturers must conform to disclosure requirements of nutritional and allergy information on their products.

We are interested in the power of certification in motivating producers to improve the quality of their products. As in the persuasion literature set in motion by Kamenica and Gentzkow (2011), we focus on the role of information disclosure through certification, and shut down all other channels of influence, such as certification fees, warranty requirements, consumer rights protection laws, and product marketing regulations.

We consider a simple, stylized model of quality certification that features a monopoly producer, a representative consumer, and a regulator. The producer chooses a quality and a price of an indivisible good. The price is freely chosen and observable, whereas the quality is costly and unobservable to the consumer. Information about the quality is only available through a certification rule that assigns a rating to each quality. The consumer observes the price and the certi-

* Date: 10th October 2019. Address: School of Economics and Finance, University of St Andrews, Castlecliffe, the Scores, St Andrews KY16 9AR, UK, az48@st-andrews.ac.uk. The author expresses his special gratitude to Hisayuki Yoshimoto for valuable suggestions and Geoff Ogle (CEO, Food Standards Scotland) for providing a practical perspective. The author also thanks Anton Kolotilin, Vladyslav Nora, Kemal Ozbek, Matthew Polisson, and three anonymous referees for helpful comments. Financial support from the Economic and Social Research Council Grant ES/N01829X/1 is gratefully acknowledged.

fication rating, and buys the good whenever its expected net value exceeds the consumer’s private outside option.

The regulator’s task is to choose a certification rule that maximizes the expected consumer surplus.¹ The regulator commits to a certification rule *ex-ante*. Thus, the producer’s quality and price decisions depend on both the certification rule and the producer’s private cost of quality. Through certification, the regulator balances the tradeoff between flexibility and control. A finer certification gives the producer more flexibility to use her private information efficiently, whereas a coarser certification gives the regulator more control to limit the producer’s monopoly power and to redistribute the surplus to the consumer.

Suppose that the probability density of the producer’s type is single-peaked. This is a standard assumption, and it includes increasing and decreasing density as special cases. Then optimal certification takes a simple form of *quality assurance*. A quality assurance rule reveals the quality if and only if it meets or exceeds a specified threshold, in effect, preventing sales of products whose quality is below the threshold.

In addition, suppose that the producer’s cost variation, which is unobservable to the regulator, is not too big.² This can be realistic in many applications: The regulator is more likely to be unable to observe (and condition upon) small cost differences. For example, a food standards agency may be uncertain about costs of hygiene among fast-food outlets. However, it will be aware of the cost difference between fast-food and sushi restaurants, thus being able to impose different quality standards. Where this assumption is met, *pass-fail certification* is optimal. A pass-fail rule is a widespread certification method in practice. It certifies whether or not the product satisfies a specified quality standard. Intuitively, the regulator sets a quality standard strictly above what any type of the profit-maximizing producer would prefer. Most producer types choose to meet this standard, at the expense of some low types in the tail of the distribution, who prefer to shut down. An optimal quality standard balances the tradeoff between a better quality by participating producers and a better participation rate.

We tackle the problem of optimal certification as follows. First, we show that it can be expressed as a delegation problem with a producer’s participation constraint. Then, using the equivalence between delegation and persuasion (Kolotilin and Zapechelnjuk 2019), we recast it as a monotone persuasion problem. Finally, we apply standard tools in the persuasion literature to solve it.

Note the standard method from the delegation literature (Amador and Bagwell 2013) does not address delegation problems with participation constraint, so using the method of Kolotilin and Zapechelnjuk (2019) is justified here and, moreover, it is easy to apply and brings additional insights by connecting the problem to the persuasion literature. Kolotilin and Zapechelnjuk (2019, Section 4) apply this

¹Our analysis can be easily adapted to the model with a welfare-maximizing regulator. This is further discussed in Section VI.

²This assumption is made precise in Section V.

method to solve the monopoly regulation problem. Although the intuition behind this method is similar, the certification problem is structurally different from the monopoly regulation problem, e.g., it features a two-dimensional decision (quality and price) and requires more stringent assumptions on the demand and cost functions to ensure existence of an interior solution. Whether or not it is easy to reduce the certification problem to the monopoly regulation problem, this paper has an advantage of providing any follow-up studies with a more fundamental method of solving it, rather than redirecting them to a different problem.

Related Literature.—There is a vast literature on certification and verifiable disclosure. The substantial part of this literature focuses on the phenomenon of unraveling and its limitations, where the quality is exogenous (we refer to Dranove and Jin (2010) for an overview, and to Ben-Porath, Dekel and Lipman (2017) and DeMarzo, Kremer and Skrzypacz (2019) for more recent models that involve Bayesian persuasion). In contrast, our paper studies certification of products with endogenous quality, with moral hazard being the central issue. In what follows, we restrict attention to papers that address certification under moral hazard.

A related paper is Zubrickas (2015) who studies optimal grading policy in a mathematically similar, but contextually different setting, where a teacher motivates a student to exert effort. The main difference from our paper is that, in effect, the teacher can impose grade-contingent transfers on the student, thus making it a problem of mechanism design with limited transfers (Condorelli 2012, Chakravarty and Kaplan 2013, Mylovanov and Zapechelnnyuk 2017), whereas transfers are prohibited in our model. As a result, optimal certification rules are very different: e.g., Zubrickas (2015) has pooling of high types, and we have pooling of low types.³

Another related paper is Albano and Lizzeri (2001) who study a certification problem in a similar production-trade setting. As Zubrickas (2015), this paper allows certification-contingent transfers. Another difference from our paper is that the certifier in Albano and Lizzeri (2001) is selfish and maximizes the proceeds from certification, whereas in our model the certifier acts in the interests of the consumer. Our paper complements Albano and Lizzeri (2001) as it targets a different set of applications and solves a different optimal design problem.

Further back in time, Shapiro (1986) addresses the question of the value of certification and licensing in a two-stage setting where disclosure at the early stage substitutes reputation at the late stage. Shapiro (1986) considers binary output, so certification essentially reduces to fully informative and uninformative.

Our paper contributes into an emerging strand of literature on Bayesian persuasion with moral hazard. Boleslavsky and Kim (2018) develop a general methodology of addressing this problem by extending the concavification method of

³Two other papers that study optimal grading with moral hazard are Dubey and Geanakoplos (2010) and Boleslavsky and Cotton (2015). Dubey and Geanakoplos (2010) consider this problem in the context of games of status (where payoffs depend on ordinal ranking). Boleslavsky and Cotton (2015) study a model of competition between certifiers; they assume binary quality, and thus only compare fully informative and uninformative certification.

Kamenica and Gentzkow (2011), and provide a technical characterization in two special cases, binary state and binary decision. Rodina (2017) considers a somewhat less general problem, where the agent has no private information when making an effort decision. Both papers consider general information disclosure structures and do not address our question of optimization on the space of partitional information structures. Feng and Lu (2016) and Zhang and Zhou (2016) study how information disclosure affects competitors' optimal spending in contests. Finally, Hörner and Lambert (2016) analyze and solve a dynamic information design model with moral hazard, with a long-lived agent whose effort decision is influenced by information disclosure to the market.

I. Model

There is a producer of an indivisible good, a representative consumer, and a regulator. The producer chooses a quality $x \geq 0$ and a price $p \geq 0$. The consumer decides whether to buy ($a = 1$) or not to buy ($a = 0$) the good. If the consumer buys the good, he obtains the net value of $x - p$, but he forgoes an outside option with value θ . The consumer's payoff is $a(x - p - \theta)$.

To produce the good of quality x , the producer pays the cost $c(x)/t$, where $c(x)$ is a cost function and t is the producer's private type that captures production efficiency (a higher type corresponds to a lower cost). If the good is sold, the producer receives its price p . The producer's payoff is $ap - c(x)/t$. Notice that the cost $c(x)/t$ is sunk and cannot be salvaged if the good is not sold ($a = 0$). This is in line with practice, where manufacturers commit to quality and production before trade.

The consumer's outside option θ and the producer's type t are independent random variables. Outside option θ is distributed on an interval normalized to be $[0, 1]$ with a continuous distribution F . Type t is distributed on an interval $[\underline{t}, \bar{t}]$, $\bar{t} > \underline{t} > 0$, with distribution G that admits a continuous density g . Cost function $c(x)$ is continuously differentiable and strictly increasing, and satisfies $c(0) = 0$. Distributions F and G , cost function c , and parameters α , \underline{t} , and \bar{t} are common knowledge.

The consumer does not observe quality x ; instead, he observes a *rating* produced by a certification rule. Formally, a certification rule is a right-continuous function $\pi : \mathbb{R}_+ \rightarrow \mathbb{R}$ that maps each quality x into a rating $\pi(x)$. Thus, a certification rule partitions the quality space and informs the consumer about the partition element that the quality belongs to.

The regulator designs a certification rule to maximize the consumer surplus. We will formalize the regulator's objective later, after required notations have been introduced.

The timing is as follows. First, the regulator designs and publicly announces a certification rule π . Then, the producer privately observes her type t and chooses a quality and a price of the good. Finally, the consumer observes his outside option θ , updates his belief about x conditional on the certification rating $\pi(x)$,

and decides whether to buy the good or not.

II. Interpreting Certification as Delegation

The problem that we pose requires optimization over the set of certification rules. However, assuming that the producer never makes strictly dominated decisions, we can represent it as a *delegation problem*. Instead of letting the producer choose any quality x and then publicizing a rating $\pi(x)$, the regulator directly restricts the set of quality levels available to the producer, and makes the producer's choice observable.

Consider any rating r that can be generated by a certification rule π , i.e., $r \in \cup_{x \geq 0} \pi(x)$. Let $\hat{x}_\pi(r)$ be the smallest quality that yields rating r ,⁴

$$\hat{x}_\pi(r) = \min\{x' : \pi(x') = r\}.$$

The consumer cannot distinguish between quality $\hat{x}_\pi(r)$ and any other quality that yields the same rating r . Because a greater quality is strictly costlier for the producer, $\hat{x}_\pi(r)$ strictly dominates any other quality x that yields the same rating, $\pi(x) = r$. Hence, a consumer who observes the rating r must realize that, among all qualities that could lead to this rating, the only producer's undominated choice is $\hat{x}_\pi(r)$. In other words, each publicly observable rating fully reveals a unique rational decision of the producer consistent with this rating. Thus, we obtain the following property of certification rules.

CLAIM 1: *For each certification rule π , each rating $r \in \cup_{x \geq 0} \pi(x)$ reveals a unique producer's undominated quality decision, $\hat{x}_\pi(r)$.*

Consequently, in this model of certification with moral hazard, the role of information disclosure is to restrict the set of undominated decisions of the producer, similar to the delegation models of Alonso and Matouschek (2008) and Amador and Bagwell (2013). This is different from standard Bayesian persuasion models, such as Rayo and Segal (2010), Kamenica and Gentzkow (2011), and Alonso and Câmara (2016), where the role of information disclosure is to manipulate beliefs of a receiver.

For a given certification rule π , the set of the producer's undominated quality choices is given by

$$X = \{\hat{x}_\pi(r) : r \in \cup_{x \geq 0} \pi(x)\}.$$

Note that X summarizes all that we need to know about π , provided the producer never makes strictly dominated decisions. Also note that choosing zero quality, $x = 0$, is always undominated, so $0 \in X$. The smallest set, $X = \{0\}$, corresponds to the completely uninformative rule that displays a constant rating regardless of the quality. The largest set, $X = \mathbb{R}_+$, corresponds to the fully informative rule that reveals the quality, e.g., $\pi(x) = x$. A finite set, $X = \{0, x_1, \dots, x_n\}$ such that

⁴Note that $\hat{x}_\pi(r)$ is well defined, since π is right-continuous.

$x_0 = 0 < x_1 < \dots < x_n < x_{n+1} = \infty$, corresponds to an interval disclosure rule that displays a distinct rating on each interval of quality $[x_i, x_{i+1})$, $i = 0, 1, \dots, n$.

We restrict attention to certification rules for which the set X is closed, as otherwise the producer's optimization problem may not have a solution.⁵

Denote by \mathcal{X} the set of closed, measurable subsets of \mathbb{R}_+ such that $0 \in X$.

III. Delegation Problem

We now represent our original certification problem as a delegation problem. In this problem, the regulator chooses a set $X \in \mathcal{X}$ of choices available to the producer (*delegation set*). Then, the producer privately observes her type t , and chooses a quality $x \in X$ and a price $p \geq 0$. Finally, the consumer observes p and x , as well as his outside option θ , and decides whether to buy the good ($a = 1$) or not ($a = 0$).

Let us formalize and simplify the producer's and the regulator's payoffs. As the consumer's payoff is $a(x - p - \theta)$, his optimal decision is $a = 1$ if and only if $x - p - \theta \geq 0$.⁶ Thus, the probability that the good with observed quality x is sold at price p is $\Pr[\theta \leq x - p] = F(x - p)$. So, the producer's payoff can be expressed as

$$(1) \quad u(t, x) = \max_{p \geq 0} pF(x - p) - \frac{c(x)}{t}.$$

For the delegation problem to be well defined, we need to impose restrictions on functions F and c , so that the profit-maximizing price p and quality x are unique for all types of the producer (except, perhaps, for a measure zero of types). To this end, we introduce two assumptions:

$$(A_1) \quad F(\theta) = \theta^\alpha, \quad \alpha > 0,$$

$$(A_2) \quad \frac{c'(x)}{x^\alpha} \text{ is strictly increasing; } \frac{c'(x)}{x^\alpha} > \bar{t} \text{ for some } x > 0.$$

Distribution F plays the role of the demand function. Assumption A_1 , though restrictive, allows for linear demand ($\alpha = 1$), concave demand where consumers with low outside option are more numerous ($\alpha < 1$), and convex demand where consumers with high outside option are more numerous ($\alpha > 1$). Furthermore, with Assumption A_1 the profit-maximizing price $p^*(x)$ can be explicitly derived

⁵For example, consider a rule $\pi(x) = \max\{1, x\}$, so $X = \{0\} \cup (1, \infty)$. So, consumers cannot distinguish qualities in $[0, 1]$, but they can distinguish each quality strictly greater than 1. If the producer's preferred quality is $x^* = 1$, this problem does not have a solution, because the producer would choose $x > 1$ that is arbitrarily close to 1. However, a simple tweak in the certification rule, e.g., $\hat{\pi}(x) = 0$ if $x < 1$ and $\hat{\pi}(x) = x$ if $x \geq 1$, leads to the same partition of the quality space, except that $x = 1$ is now separated, and $\hat{X} = \{0\} \cup [1, \infty)$ is now closed.

⁶A tie is a zero probability event.

for each quality x . Solving

$$\max_{p \geq 0} pF(x - p) = \max_{p \geq 0} p(x - p)^\alpha,$$

we obtain $p^*(x) = x/(1 + \alpha)$. Substituting $p^*(x)$ into the producer's payoff in (1) yields

$$u(t, x) = Kx^{1+\alpha} - \frac{c(x)}{t}, \quad \text{where } K = \frac{\alpha^\alpha}{(1 + \alpha)^{1+\alpha}}.$$

Assumption (A₂) ensures that the above $u(t, x)$ is strictly quasiconcave in x for each t . Moreover, the producer's optimal quality choice $x^*(t, X)$ from a set X ,

$$x^*(t, X) \in \arg \max_{x \in X} u(t, x),$$

exists for each t and is unique, except, possibly, for a measure zero of types.

Next, given a quality decision x , the regulator's payoff is equal to the consumer surplus,

$$v(x) = \int_0^{x-p^*(x)} (x - p^*(x) - \theta) dF(\theta) = \frac{\alpha}{1 + \alpha} Kx^{1+\alpha}.$$

Notice the conflict of interest between the parties: the regulator's payoff $v(x)$ is increasing in the quality x , whereas the producer wishes to choose the profit-maximizing quality that depends on his type.

In the delegation problem, the regulator chooses a delegation set X to maximize the expected consumer surplus,

$$(2) \quad \max_{X \in \mathcal{X}} \int_0^1 v(x^*(t, X)) dG(t).$$

IV. Quality Assurance Certification

A certification rule is called *quality assurance* if there is a threshold $\underline{x} \geq 0$ such that quality x is disclosed if $x \geq \underline{x}$ and remains undisclosed if $x < \underline{x}$.

Note that under any such rule, the consumer cannot distinguish among qualities in $[0, \underline{x})$. So, if the producer chooses not to meet the threshold, i.e., $x < \underline{x}$, then her choice must be $x = 0$. This is equivalent to shutting down the production, as the producer's payoff is zero in this case.

The delegation set that corresponds to the quality assurance rule with a threshold \underline{x} is $X = \{0\} \cup [\underline{x}, \infty)$, so the producer either shuts down ($x = 0$) or chooses a quality that meets or exceeds the threshold ($x \geq \underline{x}$).

The following result shows that a quality assurance rule is optimal whenever the density of the producer's type is single-peaked.

THEOREM 1: *Let (A₁) and (A₂) hold, and let g be quasiconcave. Then a quality assurance rule is optimal.*

To prove this theorem, we use Kolotilin and Zapechelnuyk (2019) to reinterpret our delegation problem as a linear monotone persuasion problem. Then we apply Kolotilin, Mylovanov and Zapechelnuyk (2019) to prove the optimality of upper-censorship rules, which translate into quality assurance rules in our setting. The complete proof is in Section VII.

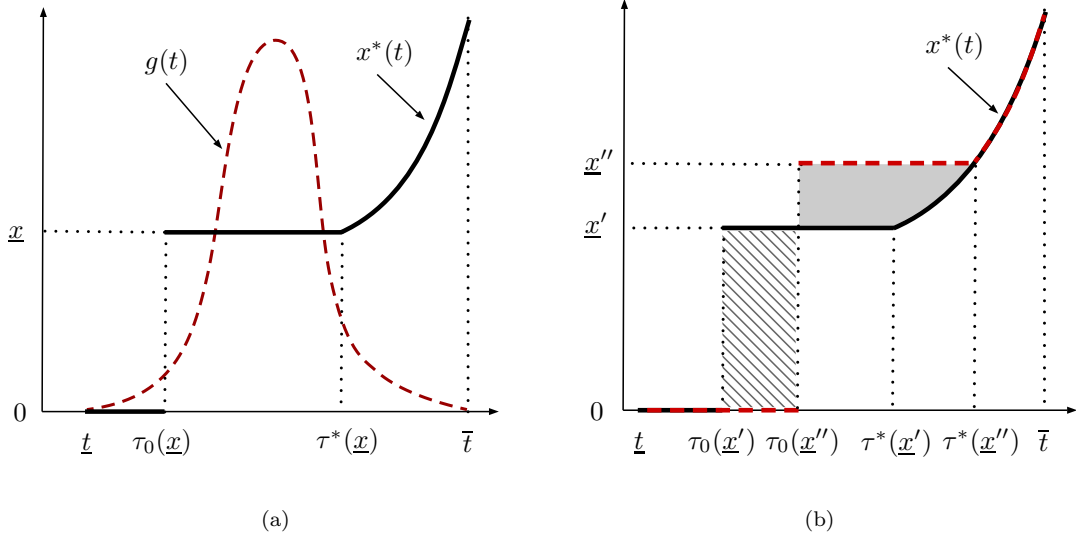


FIGURE 1. QUALITY ASSURANCE RULE.

Before discussing Theorem 1, we introduce some notations. For a given threshold $\underline{x} > 0$, let $\tau_0(\underline{x})$ be the type who is indifferent between choosing 0 and \underline{x} , and let $\tau^*(\underline{x})$ be the type for whom quality \underline{x} is the first best, so

$$u(\tau_0(\underline{x}), \underline{x}) = 0 \quad \text{and} \quad u(\tau^*(\underline{x}), \underline{x}) = \max_x u(\tau^*(\underline{x}), x).$$

Figure 1(a) illustrates a quality assurance rule with threshold \underline{x} . The set of permitted quality choices is $X = \{0\} \cup [\underline{x}, \infty)$. The solid line depicts the quality choices of different types between \underline{t} and \bar{t} , and the dashed line depicts the density g of the producer's type. The types in the left tail of the distribution, $t < \tau_0(\underline{x})$, choose quality 0. The types in the right tail of the distribution, $t > \tau^*(\underline{x})$, choose their first best quality $x^*(t)$, where

$$(3) \quad x^*(t) = \arg \max_{x \geq 0} u(t, x).$$

The remaining types that sit in the middle, $\tau_0(\underline{x}) \leq t \leq \tau^*(\underline{x})$, choose the threshold \underline{x} as their second-best choice (their first-best choice $x^*(t)$ is unavailable, as $0 < x^*(t) < \underline{x}$, so $x^*(t) \notin X$).

To understand the tradeoff that the regulator faces, let us examine what happens as a threshold rises from some x' to a greater quality x'' . The effect of this change is illustrated in Figure 1(b). The solid line and the dashed line depict the quality choices of different types before and after the change. The types between $\tau_0(x')$ and $\tau_0(x'')$ reduce their quality choice to $x = 0$ (in words, they stop producing). This is a negative effect, shown as the hatched area in Figure 1(b). However, the types between $\tau_0(x'')$ and $\tau^*(x'')$ increase their quality choice to x'' . This is a positive effect, shown as the shaded area in Figure 1(b). An optimal threshold balances these two effects.

V. Pass-Fail Certification

A certification rule is called *pass-fail* if there is a threshold $\underline{x} \geq 0$ such that the consumer is informed whether $x \geq \underline{x}$ (pass) or $x < \underline{x}$ (fail). The correspondent delegation set is $X = \{0, \underline{x}\}$. Under this rule, the producer either prefers to meet the threshold and get a “pass”, $x = \underline{x}$, or choose zero quality and get a “fail”.

Let us introduce an additional assumption:

$$(A_3) \quad u(t, x^*(t)) \geq 0 \quad \text{for all } t \in [\underline{t}, \bar{t}],$$

where $x^*(t)$ is the first-best quality for the highest type \bar{t} , given by (3). In words, any producer type gets a nonnegative payoff when choosing the quality $x^*(t)$ that is optimal for the highest type (lowest cost) producer, $t = \bar{t}$. Intuitively it means that the difference between the highest and lowest costs is small enough, so that what is best for the most efficient producer does not drive the least efficient one out of business. The main motivation for this assumption is that random variable $t \in [\underline{t}, \bar{t}]$ captures the regulator’s uncertainty about some unobservable variation in the producer’s cost. This variation is not too large, as otherwise the regulator is likely to observe it and condition the certification rule upon it. For example, a food standards agency might be unable to observe some cost differences between restaurants in the same category, but it will observe these differences between restaurants in different categories, and adjust the optimal certification to take into account the observables.

THEOREM 2: *Let (A₁), (A₂), and (A₃) hold, and let g be quasiconcave. Then a pass-fail rule is optimal.*

To understand the intuition, let us return to the quality assurance rule with threshold \underline{x} shown in Figure 1(a). Suppose that the support of types is now smaller, so $\bar{t} < \tau^*(\underline{x})$, as illustrated in Figure 2. This means that the threshold \underline{x} is greater than the producer’s first-best choice $x^*(t)$ for each type t . Even if qualities strictly above $x > \underline{x}$ were permitted, the producer would never choose them, as they are further away from his preferred choice. It is, thus, irrelevant whether choices above \underline{x} are included in the delegation set or not. So, the quality




FIGURE 2. PASS-FAIL RULE.

assurance rule with threshold \underline{x} induces the same producer's behavior as the pass-fail rule with the same threshold. Assumption (A_3) is sufficient for this to be the case when the threshold \underline{x} is chosen optimally. The formal proof is in Section VII.

VI. Discussion

The paper focuses on a simple model of certification that illustrates our solution method clearly and concisely. Let us outline the role of the focal assumptions.

The proof of Theorem 1 involves three steps: Representation of the certification problem as a delegation problem with a participation constraint, translation of the delegation problem into an equivalent monotone persuasion problem using Kolotilin and Zapechelnyuk (2019), and solution of the persuasion problem using Kolotilin, Mylovanov and Zapechelnyuk (2019). Let us discuss crucial assumptions behind each of these steps.

The first step relies on the certifier's ability to perfectly verify the producer's quality choice. In fact, we can consider a more general model where the producer's effort (input) stochastically translates into quality (output). What is important for this step is that *input* is observable and certifiable. The randomness of the output can be captured by the probability of sale at a given price, the distribution F , thus corresponding to our basic model. Certification of input is relevant in many applications. For example, a food standards agency inspects a restaurant's inputs (quality and grade of the equipment, training of the staff, cleanliness of the premises, etc.), whereas the output (the likelihood of food poisoning) is stochastic.

There are applications where certification of output is more relevant, for example, car crash tests. In this case, the producer's effort stochastically translates

into output, so the producer cannot guarantee to achieve a certain certification rating. Though we do not carry out the formal analysis of this extension of our model, conceptually this is similar. It can be addressed using the same tools and is expected to lead to qualitatively similar results.

The first step also relies on the assumption that certification deterministically maps quality to rating. This is motivated by constraints faced by certification designers in practice. For example, if the rating assignment is deliberately randomized, it is difficult to verify if the certifier followed the prescribed randomization protocol, or if he made an intentional choice in exchange of a bribe from an interested party. Furthermore, solving a model with general information disclosure rules, as in Boleslavsky and Kim (2018), can be more technically challenging. Whether the methodology of Boleslavsky and Kim (2018) can be adapted to solve our certification problem remains an open question.

In the second step, once we have a delegation problem with a participation constraint (that is, the producer can always “opt out” by choosing $x = 0$), we convert it into a balanced delegation problem (in which the extreme decisions are included in the delegation set) to apply the equivalence result of Kolotilin and Zapechelnyuk (2019). This conversion relies on Assumption (A₂), specifically, that the producer’s payoff eventually becomes negative as her quality decision increases. In essence, this is the condition of diminishing returns, which ensures the existence of an interior solution and which is often assumed in applications.

In the third step, we solve the monotone persuasion problem using an off-the-shelf technique of Kolotilin, Mylovanov and Zapechelnyuk (2019). This technique, as well as methods in Kolotilin et al. (2017), Kolotilin (2018), and Dworzak and Martini (2019), apply to a *linear persuasion problem* where the payoffs depend on the posterior mean state. Our problem is linear under Assumption (A₁) that the consumer’s outside option θ has distribution $F(\theta) = \theta^\alpha$, $\alpha > 0$.⁷ Relaxing this assumption and addressing a nonlinear persuasion problem remains an avenue for future research.

In our model, the cost of production is fixed: it is independent of the probability of sale. Our solution method can be easily adjusted to a combination of fixed and variable costs. Curiously, in the model where the cost is variable, so the producer’s payoff is $p - c(x)/t$ from a unit sold, the consumer surplus is proportional to the producer surplus when the producer chooses the price optimally. That is, the regulator’s and producer’s interests are fully aligned. So, the optimal certification is full disclosure of the quality.

We assume that a representative consumer has a private outside option. This is a standard microfoundation for a downward sloping demand. Other methods of modeling demand can also be considered.

⁷If we allow for general distribution F , a different way to satisfy the linearity property is to assume that the regulator’s payoff is a convex combination of the monopolist’s payoff and the probability that the good is sold, $F(x)$. The change of variable $y = F(x)$ makes the regulator’s and monopolist’s marginal payoffs linear in $c(F^{-1}(y))$, which is what we need to satisfy the linearity of the equivalent persuasion problem (Kolotilin and Zapechelnyuk 2019).

We also assume that the regulator maximizes the consumer surplus. This is justified by the observation that certification is often motivated by consumer protection concerns. However, our analysis can be easily adapted to the maximization of any weighted average of the producer and consumer surplus. Our results (Theorems 1 and 2) remain unchanged under a slightly stronger assumption that the density of the producer's type is log-concave, rather than quasiconcave (see Kolotilin, Mylovanov and Zapechelnyuk 2019, Theorem 2).

Finally, there is a variety of interesting and practically relevant ways to enrich the model, such as multiple dimensions of quality and partly naïve consumers, that we leave for future research.

VII. Proofs

Proof of Theorem 1. First, we use the equivalence result of Kolotilin and Zapechelnyuk (2019) (thereafter, KZ) to show that our delegation problem is equivalent to a monotone persuasion problem. Then, we apply the analysis of Kolotilin, Mylovanov and Zapechelnyuk (2019) (thereafter, KMZ) to solve the monotone persuasion problem, thus solving our delegation problem by the fact of their equivalence.

As a preparatory step, we restate our problem as a delegation problem that fits the framework of KZ. Changing the variables $y = x^{1+\alpha}$ and $z = G(t)$, so that z is uniformly distributed on $[0, 1]$, the producer's and regulator's payoffs become

$$u(z, y) = Ky - \frac{c(y^{1/(1+\alpha)})}{G^{-1}(z)} \quad \text{and} \quad v(y) = \frac{\alpha K}{1 + \alpha} y.$$

Since $G^{-1}(z) \geq G^{-1}(0) = \underline{t} > 0$, multiplying $u(z, y)$ by the positive constant $G^{-1}(z)/K$ does not change the producer's objective. Similarly, multiplying $v(y)$ by the positive constant $(1 + \alpha)/(\alpha K)$ does not change the regulator's objective. We thus obtain the equivalent expressions,

$$(4) \quad \tilde{u}(z, y) = G^{-1}(z)y - \frac{c(y^{1/(1+\alpha)})}{K} \quad \text{and} \quad \tilde{v}(y) = y.$$

The marginal payoffs are

$$(5) \quad \frac{\partial \tilde{u}(z, y)}{\partial y} = G^{-1}(z) - b(y) \quad \text{and} \quad \frac{\partial \tilde{v}(y)}{\partial y} = 1,$$

where

$$b(y) = \frac{c'(y^{1/(1+\alpha)})y^{1/(1+\alpha)-1}}{K(1 + \alpha)}$$

is strictly increasing by (A₂).

Observe that, by (A₂), there exists a finite number \bar{y} such that $u(z, \bar{y}) < u(z, 0) = 0$ for all $z \in [0, 1]$. That is, the producer strictly prefers 0 to any

$y \geq \bar{y}$. Therefore, we can restrict the domain of y to the interval $[0, \bar{y}]$ and consider the following delegation problem with the regulator's and producer's payoffs given by (4). The regulator chooses a delegation set Y , a closed subset of $[0, \bar{y}]$ that contains 0 and \bar{y} . The producer observes z , and then chooses $y \in Y$ to maximize her expected payoff. Note that $y = 0$ is included in the delegation set by assumption (because $x = 0$ is included, and $y = x^{1+\alpha}$), and the inclusion of $y = \bar{y}$ does not affect the producer's decisions, as $y = 0$ is strictly preferred to $y = \bar{y}$ for all producer types. We thus have obtained a *balanced delegation problem* (KZ, Section 2.2).

Consider now the following *monotone persuasion problem* (KZ, Section 2.3). There are an agent and a principal, whose payoffs $u_P(y, z)$ and $v_P(y, z)$ depend on an agent's decision $z \in [0, 1]$ and a state $y \in [0, \bar{y}]$. The marginal payoffs w.r.t. the decision are given by

$$(6) \quad \frac{\partial u_P(y, z)}{\partial z} = b(y) - G^{-1}(z) \quad \text{and} \quad \frac{\partial v_P(z)}{\partial z} = -1.$$

The principal chooses a monotone partitional information structure that partitions the state space $[0, \bar{y}]$ into convex sets: separating elements and pooling intervals. A monotone partition is described by a set $Y \subset [0, \bar{y}]$ of boundary points of its partition elements. Note that Y necessarily includes the extreme points 0 and \bar{y} .

In this problem, state y is initially unobserved. First, the principal chooses a monotone partition Y . Then the agent observes the partition element that contains the state, and chooses a decision $z \in [0, 1]$ that maximizes his expected payoff given the posterior belief about the state. We say that a monotone partition is optimal if it maximizes the principal's expected payoff.

The next claim is immediate by KZ's definition of the equivalence and by Theorem 1 (KZ, Section 3.1).⁸

CLAIM 2: *Set $Y \subset [0, \bar{y}]$ such that $\{0, \bar{y}\} \subset Y$ is an optimal delegation set in the balanced delegation problem with marginal payoffs (5) if and only if it is an optimal monotone partition in the monotone persuasion problem with marginal payoffs (6).*

So, to find an optimal delegation set, it suffices to solve the monotone persuasion problem with marginal payoffs (6). For that purpose, we use a result from KMZ.

First, let us recover the agent's and principal's payoffs from (6) by integrating in z with the initial value condition $u_P(y, 0) = v_P(y, 0) = 0$,⁹

$$u_P(y, z) = \int_0^z (b(y) - G^{-1}(z')) dz' \quad \text{and} \quad v_P(z) = -z.$$

⁸Note that y needs to be rescaled to be in $[0, 1]$ to fit to KZ model.

⁹The initial value condition plays no role in the analysis.

Next, let us change back the variable $t = G^{-1}(z)$ and define

$$(7) \quad \omega = b(y) = b(x^{1+\alpha}) = \frac{c'(x)}{K(1+\alpha)x^{1+\alpha}}.$$

We will refer to ω as *state*. Expressing the payoffs in these variables, we have

$$\tilde{u}_P(\omega, t) = \int_t^\omega (\omega - t)g(t)dt \quad \text{and} \quad \tilde{v}_P(t) = -G(t).$$

Let m be the posterior expectation of the state ω given the message from a disclosure rule. As the agent's payoff $\tilde{u}_P(\omega, t)$ is linear in ω , his expected payoff depends on his posterior belief about the state only through the expected state m , so

$$U(m, t) = \mathbb{E}[\tilde{u}_P(\omega, t)|m] = \int_t^m (m - t)g(t)dt.$$

As g is a positive function, $U(m, t)$ is maximized w.r.t. t at $t = m$, so the agent's optimal decision is $t^*(m) = m$.

Let us now derive the principal's expected payoff $V(m)$ conditional on m when the agent behaves optimally:

$$V(m) = \mathbb{E}[\tilde{v}_P(\omega, t^*(m))|m] = -G(t^*(m)) = -G(m).$$

As we have assumed that g (the derivative of G) is quasiconcave, it follows that $-V(m)$ is S -shaped. The following result is immediate by Corollary 1 in KMZ:

CLAIM 3: *If $-V$ is S -shaped, then an optimal signal is lower-censorship.*

So, an optimal disclosure rule is a lower-censorship signal (KMZ, Section 2.2), i.e., there exists a cutoff ω^* such that $\omega > \omega^*$ is disclosed and $\omega \leq \omega^*$ is pooled. By (7) and the strict monotonicity of b by (A₂), this is the same as disclosing y if $y > b^{-1}(\omega^*)$ and pooling y for all $y \leq b^{-1}(\omega^*)$. Translating this solution back to the balanced delegation problem, we obtain an optimal delegation set $Y = \{0\} \cup [b^{-1}(\omega^*), \bar{y}]$. Recall that adding any decisions above \bar{y} makes no difference, since the producer strictly prefers 0 to any $y > \bar{y}$. Therefore, the set $\{0\} \cup [b^{-1}(\omega^*), \infty)$ is also optimal. Finally, changing back the variable $x = y^{1/(1+\alpha)}$, we obtain an optimal delegation set for the original delegation problem, $X = \{0\} \cup [x^*, \infty)$, where x^* satisfies $b((x^*)^{1+\alpha}) = \omega^*$. This is a quality assurance rule.

Proof of Theorem 2. By Theorem 1, there exists $\underline{x} \geq 0$ such that $X = \{0\} \cup [\underline{x}, \infty)$ is a solution to problem (2).

Suppose first that $\underline{x} < x^*(\bar{t})$, where $x^*(\bar{t})$ is the unconstrained optimal choice for the highest type \bar{t} given by (3). By (A₃), all producer types prefer $x = x^*(\bar{t})$ to $x = 0$. So, raising the threshold \underline{x} to $x^*(\bar{t})$ will induce all producers who have

previously chosen quality $x < x^*(\bar{t})$ to choose $x^*(\bar{t})$ instead, which is a strict improvement for the regulator. Consequently, $x < x^*(\bar{t})$ cannot be optimal.

We thus conclude that $x \geq x^*(\bar{t})$. Observe that the producer's preferred quality $x^*(t)$ is increasing in the type, so $x^*(t) \leq x^*(\bar{t}) \leq x$ for all $t \in [t, \bar{t}]$. As the producer's payoff is strictly quasi-concave in x by (A₁) and (A₂), each producer type strictly prefers x to any quality $x > x$, as x is closer to her preferred choice. It follows that the quality assurance rule with threshold $x \geq x^*(\bar{t})$ leads to identical choices as the pass-fail rule with the same threshold x .

REFERENCES

- Albano, Gian Luigi, and Alessandro Lizzeri.** 2001. "Strategic Certification and Provision of Quality." *International Economic Review*, 42: 267–283.
- Alonso, Ricardo, and Niko Matouschek.** 2008. "Optimal Delegation." *Review of Economic Studies*, 75: 259–293.
- Alonso, Ricardo, and Odilon Câmara.** 2016. "Persuading Voters." *American Economic Review*, 106: 3590–3605.
- Amador, Manuel, and Kyle Bagwell.** 2013. "The Theory of Optimal Delegation with an Application to Tariff Caps." *Econometrica*, 81(4): 1541–1599.
- Ben-Porath, Elchanan, Eddie Dekel, and Barton L Lipman.** 2017. "Disclosure and Choice." *Review of Economic Studies*, 85: 1471–1501.
- Boleslavsky, Raphael, and Christopher Cotton.** 2015. "Grading Standards and Education Quality." *American Economic Journal: Microeconomics*, 7: 248–279.
- Boleslavsky, Raphael, and Kyungmin Kim.** 2018. "Bayesian Persuasion and Moral Hazard." Mimeo, available at <https://dx.doi.org/10.2139/ssrn.2913669>.
- Chakravarty, Surajeet, and Todd R. Kaplan.** 2013. "Optimal Allocation Without Transfer Payments." *Games and Economic Behavior*, 77: 1–20.
- Condorelli, Daniele.** 2012. "What Money Can't Buy: Efficient Mechanism Design with Costly Signals." *Games and Economic Behavior*, 75: 613–624.
- DeMarzo, Peter M, Ilan Kremer, and Andrzej Skrzypacz.** 2019. "Test Design and Minimum Standards." *American Economic Review*, 109: 2173–2207.
- Dranove, David, and Ginger Zhe Jin.** 2010. "Quality Disclosure and Certification: Theory and Practice." *Journal of Economic Literature*, 48: 935–63.
- Dubey, Pradeep, and John Geanakoplos.** 2010. "Grading Exams: 100, 99, 98,... or A, B, C?" *Games and Economic Behavior*, 69(1): 72–94.
- Dworzak, Piotr, and Giorgio Martini.** 2019. "The Simple Economics of Optimal Persuasion." *Journal of Political Economy*, forthcoming.
- Feng, Xin, and Jingfeng Lu.** 2016. "The Optimal Disclosure Policy in Contests with Stochastic Entry: A Bayesian Persuasion Perspective." *Economics Letters*, 147: 103–107.

- Hörner, Johannes, and Nicolas S Lambert.** 2016. “Motivational Ratings.” Cowles Foundation Discussion Paper No. 2035, available at <https://dx.doi.org/10.2139/ssrn.2761405>.
- Kamenica, Emir, and Matthew Gentzkow.** 2011. “Bayesian Persuasion.” *American Economic Review*, 101: 2590–2615.
- Kolotilin, Anton.** 2018. “Optimal Information Disclosure: A Linear Programming Approach.” *Theoretical Economics*, 13: 607–636.
- Kolotilin, Anton, and Andriy Zapechelnyuk.** 2019. “Persuasion Meets Delegation.” Mimeo, available at <https://arxiv.org/abs/1902.02628>.
- Kolotilin, Anton, Tymofiy Mylovanov, and Andriy Zapechelnyuk.** 2019. “Optimal Persuasion with an Application to Media Censorship.” Discussion Paper 1903, School of Economics and Finance, University of St Andrews, available at <http://ideas.repec.org/p/san/wpecon/1903.html>.
- Kolotilin, Anton, Tymofiy Mylovanov, Andriy Zapechelnyuk, and Ming Li.** 2017. “Persuasion of a Privately Informed Receiver.” *Econometrica*, 85: 1949–1964.
- Mylovanov, Tymofiy, and Andriy Zapechelnyuk.** 2017. “Optimal Allocation with Ex Post Verification and Limited Penalties.” *American Economic Review*, 107: 2666–2694.
- Rayo, Luis, and Ilya Segal.** 2010. “Optimal Information Disclosure.” *Journal of Political Economy*, 118: 949 – 987.
- Rodina, David.** 2017. “Information Design and Career Concerns.” Mimeo, available at https://cpb-us-e1.wpmucdn.com/sites.northwestern.edu/dist/d/1398/files/2016/09/Rodina_JMP-2i1a9uo.pdf.
- Shapiro, Carl.** 1986. “Investment, Moral Hazard, and Occupational Licensing.” *Review of Economic Studies*, 53: 843–862.
- Zhang, Jun, and Junjie Zhou.** 2016. “Information Disclosure in Contests: A Bayesian Persuasion Approach.” *The Economic Journal*, 126: 2197–2217.
- Zubrickas, Robertas.** 2015. “Optimal Grading.” *International Economic Review*, 56: 751–776.