(Non-)transparent (non-)mandatory tests with endogenous accuracy

Frank Rosar[*] Elisabeth Schulte[†]

May 27, 2011

1 Introduction

When information about an agent's type—the quality of his product or his ability—is important for society, there is often a discussion about whether information should be provided through a mandatory testing (or certification) procedure and whether the procedure should be transparent, i.e., whether test results should be publicly announced (cf. Faure-Grimaud et al., 2009, and Farhi et al., 2010). We tackle these questions in a setting in which the accuracy of the test is endogenous, the agent possesses imprecise private information about his quality and the agent is information-averse whereas the principal is information-loving.[1]

Faure-Grimaud et al. (2009) identify conditions under which the option to conceal the certification outcome in a (potentially) non-transparent regime is actually executed with a positive probability. They show that the value of the option to conceal is the higher for the agent, the higher the agent's ex ante uncertainty about his quality. Farhi et al. (2010) show in a setting in which certifiers pursue minimum standard certification, that the option to conceal the certification outcome is harmful for the agent, who cannot prove that he did not apply for certification. To the best of our knowledge, the issue of mandatory versus voluntary participation in testing procedures has not yet been addressed in the literature explicitly. In the literature on certification (e.g., Lizzeri, 1999), an intermediary chooses a certification fee so as to induce a level of participation which maximizes the value from certification for the agent. Instead, we take the perspective of the party who is interested in learning the quality of the agent (the principal), and ask whether it can be beneficial for her to leave the participation and/or disclosure choice at the agent's discretion. More precisely, we allow the principal to choose the accuracy of a test, i.e., the probability distribution of test results conditional on the agent's true quality, together with a transparency and a participation regime as shown in Table 1.

[*]University of Bonn. Correspondence address: email@frankrosar.de

[†]University of Mannheim. Correspondence address: elisabeth.schulte@uni-mannheim.de

[1]The principal's preference for information arises as a fairly standard property (see footnote 3). The agent's information aversion may arise as follows: Suppose the agent's income is proportional to the quality as perceived by the principal and that he is risk-averse regarding his monetary income. Then he is information-averse regarding the information the principal gathers about his quality.

	test results transparent	test results opaque
participation mandatory	MT	MO
participation voluntary	VT	VO

Table 1: Possible regimes of transparency and participation.

If a perfectly accurate test is available and neither the testing nor the disclosure of the result requires the agent's agreement, perfect information can be generated. The principal cannot achieve a superior information structure by granting the agent the right to decide upon participation and/or disclosure. Thus, if a perfectly accurate test is available, it is optimal for the principal to choose regime MT (see Table 1). However, in many applications, a perfectly accurate test is not available.

Rosar and Schulte (2010) derive the optimal testing procedure in regime VT. In this regime, optimal tests are often imperfect. The reason is that imperfections in the technology foster participation of the information-averse agent. If the test is too accurate, an information-averse individual has an incentive to take the test only if his signaling motive is strong enough, i.e., if he has observed a sufficiently favorable private signal about his quality. The worst case inference in the case of non-participation is limited to the lowest realization of private information. If the worst case inference is that the agent is bad for sure, an unraveling argument applies. Adverse beliefs render participation in any test preferable to non-participation. If the agent's private information is imperfect in the sense that he cannot know to be bad for sure, the worst case inference can be better for him than participating in a sufficiently accurate test. The imperfectness of the private information in this sense "insures" the agent against unraveling. In order to motivate an agent with a low realization of private information to take the test, the test has to allow to some extend for pooling. If the test issues imperfect signals, good and bad quality are not perfectly identifiable as such. This allows an agent who is more likely to be bad to pool (on average) with better quality. Technical constraints limiting the accuracy of the testing procedure may not be binding in this regime. While the introduction of technical constraints allows for a clearly worse inference in the MT regime, the inference in the VT regime may not be affected, rendering this regime relatively better.

In the present paper, we show that regimes MT, MO and VO are equivalent for the principal such that the regime choice problem can be solved by comparing voluntary with mandatory participation in a regime with transparent test results. We show that if a perfectly accurate test is not available, voluntary participation may generate better information about the agent's quality than mandatory participation. With voluntary participation, the principal may be able to deduce useful information from the agent's participation behavior. For this channel to be valuable, the agent's private signal has to be sufficiently informative. On the other hand, for the channel to be available, the agent's private signal must not be too informative. Moreover, the maximum feasible accuracy of the test must neither be too high nor too low.

2 The model

We consider a principal–agent framework with three distinct features: (1) The principal is interested in the agent's quality and the agent is interested in how his quality is assessed by the principal, (2) the agent possesses only a probabilistic assessment of his quality, and (3) the principal may learn something about the agent's quality through a test. We are interested in how different testing regimes affect equilibrium learning when the accuracy of tests is constrained. Testing regimes differ in whether participation in the test is mandatory (M) or voluntary (V), and whether testing results are transparent (T) or opaque (O).

The agent's quality is either good ($\omega = 1$) or bad ($\omega = 0$). The agent observes only a private signal s, drawn from a finite set $\mathcal{S} = \{1, \ldots, S\}$, correlated with ω. The joint distribution of s and ω is described by the marginal probability of private signal s, $p_s \in (0, 1]$ with $\sum_s p_s = 1$, and the conditional probability that the agent's quality is good, given that his private signal is s, θ_s. Let $0 \leq \theta_1 < \theta_2 < \ldots < \theta_S \leq 1$ and let $\sum_s p_s\theta_s \in (0, 1)$. When the agent takes the test, he either passes (L) or fails (H). A test is characterized by the conditional probability of passing the test given the agent's quality. For simplicity, we consider tests which are only subject to one type of error, false positives.[2] A good agent always passes the test, and a bad agent passes with probability $\varepsilon \in [0, 1)$.

The timing is as follows: First, the principal chooses a testing procedure $(\varepsilon, \rho, \tau) \in [\underline{\varepsilon}, 1) \times \{M, V\} \times \{T, O\}$ and the agent observes the principal's choice. $\underline{\varepsilon} \in [0, 1)$ describes the technological constraint which limits the maximum accuracy of the test. Second, nature draws the agent's type ω and his private signal s. Third, the agent takes the test or not. The vector $\alpha = (\alpha_1, \ldots, \alpha_S) \in [0, 1]^S$ describes the probabilities with which the agent participates when his private signal is s. To describe the agent's behavior, we introduce the vector notation $\mathbf{1} = (1, \ldots, 1)$ and $\mathbf{0} = (0, \ldots, 0)$. If $\rho = V$, the agent decides whether to participate or not. If $\rho = M$, the agent is forced to participate. I.e., $\alpha = \mathbf{1}$. Fourth, when the test is used, nature draws a test result and sends it to the agent. Fifth, the test result might be revealed to the principal. The vector $\beta^\sigma = (\beta_1^\sigma, \ldots, \beta_S^\sigma) \in [0, 1]^S$, $\sigma \in \{L, H\}$, describes the probabilities with which the test result is revealed to the principal given that the agent has observed signal s and test result σ. If $\tau = T$, any test result $\sigma \in \{L, H\}$ that is generated is automatically revealed. I.e., $\beta^L = \beta^H = \mathbf{1}$. If the principal observes nothing ($\sigma = \emptyset$), she can infer that the agent did not take the test. If $\tau = O$, it depends on the agent's disclosing behavior what the principal observes. A test result is verifiable, but not having taken the test is not verifiable. The principal does not observe a test result ($\sigma = \emptyset$) if the agent either did not take the test or if the agent does not reveal the test result. Otherwise, she observes the true test result $\sigma \in \{L, H\}$. Sixth, upon the observation of $\sigma \in \{L, H, \emptyset\}$, the principal updates her assessment of the joint distribution of the agent's private signal and his true quality. She makes her subsequent decisions dependent on this assessment. In the first part of the paper we pursue a reduced form modeling approach where the consequences of subsequent decisions are summarized in reduced form utility functions which directly depend on the probability with which she believes the agent to be good, which we call quality belief μ_σ. μ_σ is a function of the belief in the

[2] Tests with this type of error are, e.g., common in software testing. If a software application contains no bugs, no bugs are detected. If it contains a bug, it is not always detected. It follows from Rosar and Schulte (2010) that our argument carries over to the case where the tests are subject to both types of error.

classical game–theoretical sense associated to the information set σ.

Reduced–form payoffs are as follows: The agent wants to be perceived as good and is risk-averse with respect to how he is perceived. He evaluates the induced quality beliefs with the strictly increasing and strictly concave utility function $u : [0, 1] \to \mathbf{R}$. We normalize $u(0) = 0$. The principal's utility function $v : [0, 1] \to \mathbf{R}$ is convex in μ_σ, capturing information-loving preferences.[3]

Our equilibrium concept is an adaptation of sequential equilibrium to the reduced-form setting. An equilibrium of a testing procedure $(\varepsilon, \rho, \tau)$ is characterized by a participation behavior α (if $\rho = V$), a disclosing behavior β^L and β^H (if $\tau = O$), and a system of quality beliefs μ_L, μ_H and $\mu_\emptyset$ such that (i) the agent's behavior is sequentially rational given the quality beliefs, (ii) there exists a sequence of strictly mixed behavior[4] converging to the equilibrium behavior such that the quality beliefs associated with the elements of the sequence converge towards the equilibrium quality beliefs, and (iii) quality beliefs are derived from Bayes' Law whenever possible.

The principal chooses a testing procedure and an equilibrium so as to maximize her ex ante expected utility.

3 Analysis

Quality beliefs and optimization problems

For future reference, we derive the principal's quality beliefs after the observation of $\sigma \in \{L, H, \emptyset\}$ for the case where Bayes' Law is applicable. In this case, quality beliefs are determined by the agent's supposed behavior. When Bayes' Law is applicable, i.e., when the denominator of the following expressions is positive, quality beliefs obey the following functional form:

$$\mu_\sigma(\alpha, \beta^\sigma) = \frac{\sum_s p_s \alpha_s \beta_s^\sigma \theta_s \text{Prob}(\sigma|\omega = 1)}{\sum_s p_s \alpha_s \beta_s^\sigma (\theta_s \text{Prob}(\sigma|\omega = 1) + (1 - \theta_s)\text{Prob}(\sigma|\omega = 0))} \tag{1}$$

for $\sigma \in \{L, H\}$ and

$$\mu_\emptyset(\alpha, \beta^L, \beta^H) = \frac{\sum_s p_s (1 - \alpha_s \beta_s^H)\theta_s}{\sum_s p_s ((1 - \alpha_s \beta_s^H)\theta_s + (1 - \alpha_s(\beta_s^H \varepsilon + \beta_s^L (1 - \varepsilon)))(1 - \theta_s))}. \tag{2}$$

When Bayes' Law is not applicable, quality beliefs are the limits of sequences of quality beliefs which obey the functional form in (1) and (2).

The agent takes the quality beliefs as given. His actions induce a (potentially degenerate) lottery over these quality beliefs. He chooses his participation and disclosure actions so as to maximize his ad interim expected utility $\sum_\sigma \text{Prob}(\sigma|s, \text{A's actions}) u(\mu_\sigma)$.

The principal chooses a testing procedure and an associated equilibrium so that the induced quality belief lottery maximizes her ex ante expected utility $\sum_{s,\sigma} p_s \text{Prob}(\sigma|s, \text{A's eq. behavior}) v(\mu_\sigma)$. If there is

[3]For a non-reduced example, consider a setting in which the principal makes a choice from a set of actions and where each action gives rise to a state-dependent payoff for the principal (i.e., a payoff dependent on the agent's quality). Her expected payoff from a given action is linear in the probability μ_σ that the agent's quality is good. The principal's reduced-form payoff function is the upper envelope of a set of linear functions, which is (weakly) convex in μ_σ. See also subsection 4.

[4]Our equilibrium concept requires strict mixing only for decisions that the agent actually makes, which differs across regimes.

no learning in equilibrium, the quality belief is the prior probability that the agent is good, $\mu_0 = \sum_s p_s \theta_s$, with probability one. If there is learning, the quality belief lottery is a mean preserving–spread of the prior quality belief. The convexity of the principal's utility function and Jensen's Inequality imply that she (weakly) prefers the mean–preserving spread to the prior. Hence, the convexity of v can be interpreted as a preference for learning. More generally, we have:

Lemma 1 *The principal (weakly) prefers one quality belief lottery over another if the former is a mean–preserving spread of the latter.*

Corollary 1 *Let $\underline{\mu} \leq \overline{\mu}$. The principal (weakly) prefers a quality belief lottery with support $\{\underline{\mu}, \overline{\mu}\}$ and expected quality belief μ_0 over any quality belief lottery with the same expected quality belief where the support is a subset of $[\underline{\mu}, \overline{\mu}]$.*

Equivalence of testing regimes MT, MO and VO

Our first result is that testing regimes MT, MO and VO are equivalent from a design point of view. Learning about the agent's type can occur through two channels: directly through information generated by the test and indirectly through private information. The result is established by showing that optimal learning only through the first channel can be obtained in all three regimes and that learning through the second channel is not possible in any of the regimes.[5]

Consider a test with a fixed accuracy. Optimal learning only through generated information is achieved when the agent always takes the test and the test result is always revealed. This is what happens in regime MT. In regime MO, perfect inference of the test result follows from sceptical beliefs and an unraveling argument,[6] participation is not an issue as it is mandatory.[7] In regime VO, sceptical beliefs lead to the same unraveling argument as in regime MO, given that the agent always participates. Since the testing procedure is opaque, the agent can conceal any unfavorable test result such that there are no cost associated to participation. It follows that participation is optimal and a full participation equilibrium exists.

It remains to argue that learning through private information is not possible for the principal. Learning through private information requires that the agent makes one of his decisions in a non–trivial way dependent on his private information. After the agent obtains a certain test result, his utility is not affected by his private information. Therefore the revelation decision might only depend on the agent's information in a non–trivial way if revelation and non–revelation are associated with the same quality belief. Although the agent's private information affects his expected utility from participation, participation has an option character such that his expected utility from participation cannot be lower than that from non–participation. Therefore the participation decision might only depend on the agent's private information in a non–trivial way if the option value is zero inducing a degenerate quality belief

[5] Interestingly, inference about the agent's private information may be possible in equilibrium, while learning about the agent's type through private information is not. See Footnote 8 below.

[6] There exists a continuum of equilibria with perfect inference of the test result. Any $\beta^L \in [0,1]^S$ may be part of such an equilibrium. Non–revelation of a test result and revelation of L then both allow for perfect inference that the agent failed.

[7] A similar point is made in Section 5 in Faure-Grimaud et al. (2009). There, if hiring a certifier is observable, then the option to conceal a certification outcome is worthless.

lottery. It follows that even if the agent's behavior depends on his private information in a non–trivial way, the induced quality belief lottery does not. Hence, learning about the agent's type through private information is not possible.[8]

Proposition 1 *Consider a test with accuracy* ε*. The for the principal best equilibrium in regimes MT, MO and VO induces the same lottery over quality beliefs.*

Three conclusions follow: Firstly, voluntary participation may only improve the principal's information as compared to mandatory participation if the testing procedure is transparent. Secondly, we can simplify the principal's optimization problem to optimally choosing the accuracy of the test in either regime MT or VT. If VT beats MT, VT beats all the other regimes, too. Thirdly, our comparison of transparent regimes with mandatory versus voluntary participation implies a comparison of mandatory versus voluntary disclosure in regimes with voluntary participation. Hence, if the principal prefers a test in regime MT to the tests in regime VT, but it is not feasible to force the agent to participate, she can instead grant the agent the right to decide upon disclosure.

Superiority of VT requires an informative, imperfect private signal

Consider first the two benchmark cases in which the agent either has no informational advantage or he is perfectly informed. In neither case is regime VT strictly preferred over regime MT.

If the agent's private information is uninformative, his expected quality is correctly inferred when he does not take the test. When he takes the test, he incurs a fair, non–degenerate quality belief lottery. As he is risk–averse, he strictly prefers not to take the test. Neither learning through generated information nor learning through private information is feasible in regime VT. Hence, an informative private signal—and therewith the existence of a signaling motive on the agent's side—is necessary for a superiority of regime VT.

If the agent is perfectly informed about his quality, adverse beliefs render full participation optimal and induce the same quality belief lottery as in regime MT. When the test is perfect, it follows that perfect learning of the agent's quality is possible in both regimes. To establish that a superiority of VT is impossible, we argue that learning through private information is impossible with an imperfect test. Since the agent cannot at the same time be indifferent between taking and not taking an informative test when he is good and when he is bad, there exists at most one action which is taken only by one type of agent. Since the agent has a higher incentive to take an informative test when he is good than when he is bad, we have to distinguish only two cases. First, suppose that only the bad agent does not participate. Then the principal's belief associated to non–participation is that he is bad such that participation is strictly optimal as it allows the bad agent to pool with a positive probability. Second,

[8] When the testing procedure is opaque and the accuracy of the test is not perfect, there can exist equilibria where the agent's behavior depends in a non–trivial way on his private information. One equilibrium type is the following: The agent always take the test, but he reveals a positive test result only if his private information is sufficiently low. The revelation of H conveys positive information (good test result) and negative information (low private signal) which cancel out in the updating process. The induced quality belief distribution is degenerate (no learning through private information), but the observation of σ allows for non–trivial inferences about the agent's private information (learning about private information).

suppose that only the good agent participates. Then the belief associated to passing is that the agent is good. As the agent passes with probability one when he is good, he has a strict incentive to take the test. However, then the belief associated to non–participation is that the agent is bad, which yields a strict incentive to participate for the agent also if he is bad. It follows that an imperfection of the agent's private information is necessary for the superiority of regime VT.

Lemma 2 *Consider a test with accuracy ε. (a) If the agent's private signal is uninformative, there exists a unique equilibrium in regime VT which induces a degenerate quality belief lottery. (b) If the agent's private signal is perfect, the best equilibrium in regime VT induces the same quality belief lottery as regime MT.*

Corollary 2 *If the agent's private signal is either uninformative or perfect, the principal (at least weakly) prefers regime MT over regime VT.*

To avoid case distinctions in the further analysis we henceforth assume:

Assumption 1 $S > 1, \theta_1 > 0$.

Superiority of VT requires a sufficiently high maximum test accuracy

For a test of given accuracy, there is no difference between regime MT and regime VT if the agent in equilibrium endogenously always participates. A potential advantage of regime VT prevails only if it is possible to infer private information from the observed participation behavior. However, if ε is too high, then in any equilibrium in which learning is possible, i.e., in which the quality belief lottery is non-degenerate, there is no separation of private information. The higher ε, the less the quality belief distribution associated with taking the test depends on the private signal. The quality belief associated with not taking the test does not depend on the private signal either. Hence, the lower the accuracy of the test, the less the expected utility form either action differs for different private signals. As a consequence, there are only equilibria with pooling of private information.

Lemma 3 *There exists $\underline{\varepsilon}'$, such that if $\underline{\varepsilon} \geq \underline{\varepsilon}'$, the principal is indifferent between regime MT and any equilibrium in which learning is possible in regime VT.*

If $\underline{\varepsilon}$ is too high, even with the lowest realization of private information, the agent is willing to voluntarily participate at all feasible tests in regime VT. In order to induce separation of private information, the test has to be sufficiently accurate. The more accurate the test is, the less is bad quality pooled with good quality. This in turn makes participation less attractive for the agent if his private signal is low and eventually deters participation.

Superiority of VT requires a sufficiently low maximum test accuracy

If the accuracy of the test is sufficiently high, the lotteries over the induced quality beliefs differ across regimes. In regime VT, the agent does not take the test if his private signal is too low. Therewith, passing a test with voluntary participation becomes a more favorable signal than in the case of mandatory participation, $\mu_H^{VT} > \mu_H^{MT}$. As a downside, an intermediate belief $\mu_\emptyset^{VT}$ realizes with a positive probability. Lemma 1 implies that the principal likes the former effect of endogenous participation, and Corollary 1 implies that she dislikes the latter. Without a further specification of preferences, it is not possible to rank the regimes in general.

As the test accuracy increases, the advantage of regime VT vanishes, because the difference between μ_H^{VT} and μ_H^{MT} becomes small. At the same time the disadvantage of regime VT becomes more pronounced as the agent more often refrains from taking the test. Hence, the principal has a strict preference for regime MT if the maximum test accuracy is very high.

Lemma 4 *There exists $\underline{\varepsilon}''$, such that if $\underline{\varepsilon} \leq \underline{\varepsilon}''$, the principal strictly prefers some available test in regime MT to all available tests in regime VT.*

A consequence of the previous two Lemmata is that regime VT can be superior only if the maximum test accuracy is intermediate.

Sufficient conditions for the superiority of VT

We construct a simple setting with binary private information which allows us (i) to easily show that a regime with voluntary participation can be strictly optimal, and (ii) to give an impression under which circumstances voluntary participation is likely to yield a superior information structure. Our strategy to show the superiority of a regime with voluntary participation is to compare the best possible test in regime MT (the test with maximum feasible accuracy) to a particular test (and an associated equilibrium) in regime VT, where perfect separation of private information is induced. We assume that the agent evaluates the principal's quality beliefs with the utility function $u(\mu) = \sqrt{\mu}$ and show that there exist preferences for the principal such that she prefers regime VT over regime MT. The agent's signal s is either l or h, with $0 < \theta_l < \theta_h < 1$.

In our setting with imperfect private information, there always exists a test in regime VT and an associated equilibrium which allows for perfect inference of the agent's private information.

Lemma 5 *There exists $[\varepsilon', \varepsilon''] \subset [0, 1)$ with $\varepsilon' < \varepsilon''$ such that for each $\varepsilon \in [\varepsilon', \varepsilon'']$ an equilibrium exists in which the agent participates at test (ε, V, T) if and only if $s = h$.*

Proof. Consider the testing procedure (ε, V, T) and suppose the agent participates if and only if $s = h$. In this case, the principal's quality beliefs are $\mu_\emptyset^{VT} = \theta_l$, $\mu_L^{VT} = 0$ and $\mu_H^{VT} = \theta_h/(\theta_h + (1 - \theta_h)\varepsilon)$. It is indeed optimal for an agent with private signal $s = h$ to participate if $(\theta_h + (1 - \theta_h)\varepsilon)u(\mu_H^{VT}) \geq u(\mu_\emptyset^{VT})$ which is equivalent to $\varepsilon \geq (\theta_l - \theta_h^2)/(\theta_h(1 - \theta_h))$. Let $\varepsilon' := \max\{0, (\theta_l - \theta_h^2)/(\theta_h(1 - \theta_h))\}$. If he takes the test, an agent with private signal $s = l$ obtains a payoff $(\theta_l + (1 - \theta_l)\varepsilon)u(\mu_H^{VT})$ which is strictly less than $(\theta_h + (1 - \theta_h)\varepsilon)u(\mu_H^{VT})$ and increasing in ε. Hence, there is a $\varepsilon'' > \varepsilon'$ such that it is optimal for an agent with private information $s = l$ not to participate if $\varepsilon \leq \varepsilon''$. As $\theta_l u(\mu_H^{VT}) < u(\mu_\emptyset^{VT})$ and $u(\mu_H^{VT}) > u(\mu_\emptyset^{VT})$, we have that $\varepsilon'' \in (0, 1)$. q.e.d.

Next, we show that there exist an intermediate maximum accuracy of the test and reduced-form preferences for the principal such that she strictly prefers a regime with voluntary participation over mandatory participation.

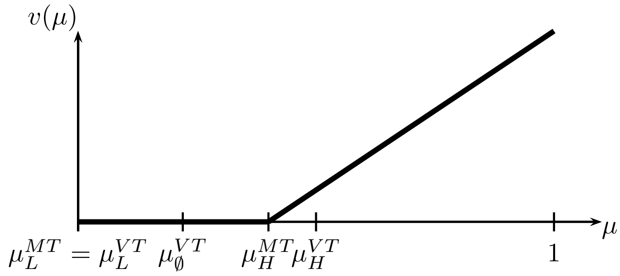

Figure 1: Example for $v(\mu)$ with $\mu' = \mu_H^{MT}$ and quality beliefs generated in regimes MT and VT, respectively. With the parameters $\theta_h = 1/3$, $p_h = 1/4$, $\theta_l = 2/9$, $\varepsilon = 1/2$, we obtain $\mu_L^{MT} = \mu_L^{VT} = 0$, $\mu_H^{MT} = 2/5$, $\mu_H^{VT} = 1/2$ and $\mu_\emptyset^{VT} = \theta_l = 2/9$.

Proposition 2 *There exists* $\underline{\varepsilon} \in (0, 1)$ *and* $v(\mu)$ *weakly convex such that the principal strictly prefers regime VT over regime MT.*

Proof. Consider $\underline{\varepsilon} = \varepsilon''$, as implicitly defined in the proof of the previous lemma. Under mandatory participation any $\varepsilon \in [\underline{\varepsilon}, 1)$ induces a lottery over the two quality beliefs $\mu_L^{MT} = 0$ and $\mu_H^{MT} = \mu_0/(\mu_0 + (1-\mu_0)\varepsilon)$. It follows from Corollary 1 that choosing $\varepsilon = \underline{\varepsilon}$ is (weakly) optimal under mandatory participation.
Compare now the optimal test under mandatory participation, (ε'', M, T), with the specific test with voluntary participation (ε'', V, T) where the agent participates if and only if $s = h$ (as was shown to be incentive compatible in the preceding lemma). The highest possible quality beliefs under mandatory and voluntary participation are then $\mu_H^{MT} = \mu_0/(\mu_0 + (1-\mu_0)\varepsilon'')$ and $\mu_H^{VT} = \theta_h/(\theta_h + (1-\theta_h)\varepsilon'')$, respectively. Since $\mu_0 < \theta_h$, $\mu_H^{MT} < \mu_H^{VT}$. Let $v(\mu) := \max\{0, \mu - \mu'\}$, where $\mu' \in [\mu_H^{MT}, \mu_H^{VT})$. Then the principal obtains zero expected utility when she forces participation while she obtains positive expected utility when she leaves participation at the agent's discretion. q.e.d.

To get an intuition for Proposition 2, consider $v(\mu) = \max\{0, \mu - \mu'\}$, with $\mu' \in (0, 1)$. Such a payoff function for the principal may stem from optimal task-assignment, where one task generates zero payoff and the other task generates payoff 1 if $\omega = 1$ and $-\mu'/(1-\mu')$ if $\omega = 0$. Then, μ' is the threshold that $\text{Prob}(\omega = 1|\sigma)$ must pass to render the assignment of the risky task optimal. If participation is voluntary and the agent participates only if his private information is high, an agent who passes the test is good with a high enough probability such that it pays to assign the risky task. If participation is mandatory instead, the informational value of the test result is diluted and the probability that the agent is good does not pass the threshold μ'. In this case, the generated information has no value for the principal, because she chooses the safe task independently of the test result.

A regime with mandatory participation also has no value if the prior is too pessimistic, i.e., if $\mu_0 \leq \mu'\varepsilon/(1 - \mu'(1-\varepsilon))$. With voluntary participation, valuable information can be generated if the high private signal is sufficiently informative, i.e., if $\theta_h > \mu'\varepsilon/(1-\mu'(1-\varepsilon))$, and the test is intermediately accurate (such that an equilibrium with separation of private information exists). Our analysis suggests that regime VT is superior if the prior quality expectations are low.

4 A non-reduced example

We now demonstrate an example for a non-reduced game in which the individuals' preferences are defined over money. Our example leads to the reduced-form preferences as considered in the preceding subsections and voluntary participation yields superior information. After the testing stage, the principal and the agent enter the following game.

The principal chooses the scale of a project, characterized by an investment $I \in [0, k-1]$ with $k > 1$ which yields a monetary payoff $Y(I) := k(I/(k-1))^{(k-1)/k}$ if the project is successful and zero else. The project succeeds if and only if the agent's quality is high, i.e., if $\omega = 1$. The expected project surplus from investment choice I when the agent's quality is good with probability μ is thus $-I + \mu Y(I)$. By choosing the investment which maximizes the expected surplus, the generated expected project surplus is $\pi(\mu) = \mu^k$. We assume that the principal is risk-neutral such that his utility from monetary payoff t_P is $\tilde{v}(t_P) = t_P$.

Principal and agent engage in bargaining. We assume that Y is either not verifiable or it is observable only after a long delay, such that the agent obtains a fixed transfer. Bargaining yields a fraction δ of the expected surplus for the agent and an expected monetary payoff in hight of a fraction $1-\delta$ of the expected surplus for the principal.[9]

Hence, the principal's choice of the project scale maximizes the expected project surplus, inducing a reduced form utility function $v(\mu) = (1-\delta)\mu^k$. The (state-independent) monetary transfer payed to the agent is $t_A = \delta\pi(\mu)$. The agent is risk-averse. His preferences are captured by the utility function $\tilde{u}(t_A) = (t_A/\delta)^{1/(2k)}$, yielding a reduced-form utility function $u(\mu) = \sqrt{\mu}$, as in the preceding subsection.

For the same parameters as in the example depicted in Figure 1, we obtain the same participation behavior and the lotteries over quality beliefs induced by the testing procedures $(1/2, M, T)$ and $(1/2, V, T)$, respectively, are also the same. With mandatory participation, the principal's expected payoff is $5/8(2/5)^k$. With voluntary participation, she obtains $3/4(2/9)^k + 1/6(1/2)^k$. Leaving the participation choice at the agent's discretion is optimal for the principal if $k \geq 6$.

In order to gain additional insights into the conditions that yield a superiority of regime VT, we perform the following comparative statics exercise. Suppose that the model parameters are such that an equilibrium with separation of private information for $\varepsilon = \underline{\varepsilon}$ exists in regime VT. Suppose further that such an equilibrium still exists after the proposed changes of the parameters. A sufficient condition for the principal to strictly prefer regime VT is $p_h\theta_h \left(\frac{\theta_h}{\theta_h + (1-\theta_h)\underline{\varepsilon}}\right)^k > \mu_0 \left(\frac{\mu_0}{\mu_0 + (1-\mu_0)\underline{\varepsilon}}\right)^k$, which is equivalent to:

$$\frac{p_h\theta_h}{p_l\theta_l} > \frac{\left(\frac{\mu_0}{\mu_0+(1-\mu_0)\underline{\varepsilon}}\right)^k}{\left(\frac{\theta_h}{\theta_h+(1-\theta_h)\underline{\varepsilon}}\right)^k - \left(\frac{\mu_0}{\mu_0+(1-\mu_0)\underline{\varepsilon}}\right)^k}$$

As the RHS of the above inequality is strictly increasing in k and the LHS is not affected by a change in k, regime VT is superior if k is high enough. Next, note that the LHS is strictly decreasing in θ_l. All else equal, an increase in θ_l is accompanied by an increase in μ_0. The RHS is strictly increasing in μ_0. Consequently, if regime VT is superior if $\theta_l = \theta_l'$, then it is superior if $\theta_l < \theta_l'$ (provided the existence of

[9]Note that the impossibility to condition the agent's payment on Y precludes transmission of information in the bargaining stage.

a separating equilibrium which requires that θ_l does not fall below a certain bound). Similarly, consider an increase in θ_h accompanied by a change of θ_l such that μ_0 stays constant and $\theta_h - \theta_l$ increases. The LHS increases, and the RHS decreases. Thus, regime VT is also superior if the private signal is sufficiently informative. We conclude that regime VT is the more likely to be superior (i) the more convex the principal's payoff function, (ii) the lower the probability that the agent's quality is good if he has observed the low private signal (up to the point where the equilibrium with separation of private information ceases to exist), (iii) and the more informative the agent's private signal (with the same qualification).

5 Conclusion

In an environment in which the reliable identification of good quality is very important, and in which a test can generate information, but is subject to false positives, a transparent testing procedure with voluntary participation can generate better information than a testing procedure with mandatory participation or a testing procedure with an opaque disclosure regime.

In a transparent regime, voluntary participation has the advantage that it may allow access to the agent's private information. If the participation behavior depends on the agent's private information, then in the case that the agent passes the test, the probability that his quality is in fact good is higher than in the case of mandatory participation. The disadvantage of voluntary participation is that test results are observed with a smaller probability. An intermediate quality belief realizes with a positive probability, and less probability mass is on the extreme quality beliefs. The advantage of voluntary participation outweighs its disadvantage if the principal's utility function is convex enough or if it is important that the quality belief passes a certain threshold. Further, our analysis suggests that a transparent regime with voluntary participation is superior in environments in which the principal's prior quality expectations are low, or the agent's interim quality expectations are sufficiently diverse and at the same time not too extreme.

In our argument, we have abstracted from the use of participation fees. If the principal can offer a compensation for participation, she does not have to compromise on the accuracy of the test for a higher probability of participation. With an appropriately chosen fee (which may be negative), she can achieve the same information structure in a regime with voluntary participation as in the regime with mandatory participation, or even one that she prefers. Our analysis suggests that if the principal's stakes in the decision problem are sufficiently higher than they are for the agent, instead of forcing the agent to take a particular (imperfect) test, the principal (i) prefers to offer him a payment for taking the test if the maximum accuracy of the test is high, and (ii) charge a participation fee if the maximum accuracy of the test is low.

Appendix

Proof of Lemma 1

Let $\widetilde{\mu}_1$ and $\widetilde{\mu}_2$ be random variables where $\widetilde{\mu}_2$ is a mean–preserving spread of $\widetilde{\mu}_1$. I.e., there exists a random variable $\widetilde{\mu}'$ with $\mathbf{E}[\widetilde{\mu}'|\widetilde{\mu}_1] = 0$ such that $\widetilde{\mu}_2$ and $\widetilde{\mu_1} + \widetilde{\mu}'$ have the same distribution. Hence, $\mathbf{E}[v(\widetilde{\mu}_2)] = \mathbf{E}[v(\widetilde{\mu}_1 + \widetilde{\mu}')]$. By the Law of Iterated Expectations, $\mathbf{E}[v(\widetilde{\mu}_1 + \widetilde{\mu}')] = \mathbf{E}[\mathbf{E}[v(\widetilde{\mu}_1 + \widetilde{\mu}')|\widetilde{\mu}_1]]$. By Jensen's Inequality and (weak) convexity of v, $\mathbf{E}[\mathbf{E}[v(\widetilde{\mu}_1 + \widetilde{\mu}')|\widetilde{\mu}_1]] \geq \mathbf{E}[v(\mathbf{E}[\widetilde{\mu}_1 + \widetilde{\mu}'|\widetilde{\mu}_1])]$. By $\mathbf{E}[\widetilde{\mu}'|\widetilde{\mu}_1] = 0$, $\mathbf{E}[v(\mathbf{E}[\widetilde{\mu}_1 + \widetilde{\mu}'|\widetilde{\mu}_1])] = \mathbf{E}[v(\widetilde{\mu}_1)]$. Hence, $\mathbf{E}[v(\widetilde{\mu}_2)] \geq \mathbf{E}[v(\widetilde{\mu}_1)]$. q.e.d.

Proof of Corollary 1

Let $\widetilde{\mu}_2$ be the random variable with support $\{\underline{\mu}, \overline{\mu}\}$ and expected value μ_0, and let $\widetilde{\mu}_1$ be the random variable where the support is a subset of $[\underline{\mu}, \overline{\mu}]$ and which has the same expected value. Define a new random variable $\widetilde{\mu}'$ which depends on the realization μ of $\widetilde{\mu}_1$. $\widetilde{\mu}'$ assumes value $\overline{\mu} - \mu$ with probability q and value $\underline{\mu} - \mu$ with probability $1 - q$ where q is chosen such that $(1-q)(\underline{\mu} - \mu) + q(\overline{\mu} - \mu) = 0$. It follows that $\widetilde{\mu}_2$ and $\widetilde{\mu}_1 + \widetilde{\mu}'$ have the same binary support. Since $\widetilde{\mu}_2$ and $\widetilde{\mu}_1 + \widetilde{\mu}'$ also have the same expected value, they must be distributed according to the same distribution. I.e., $\widetilde{\mu}_2$ is a mean–preserving spread of $\widetilde{\mu}_1$. q.e.d.

Proof or Proposition 1

Regime MT is non–strategic. It induces a unique binary lottery over quality beliefs $\mu_H^{MT} = \mu_H(\mathbf{1}, \mathbf{1}) = \mu_0/(\mu_0 + (1 - \mu_0)\varepsilon)$ and $\mu_L^{MT} = \mu_L(\mathbf{1}, \mathbf{1}) = 0$.

Consider regime MO. Revelation behavior $\beta^H = \mathbf{1}$, $\beta^L = \mathbf{0}$ together with quality beliefs $\mu_H^{MO} = \mu_H^{MT}$ and $\mu_\emptyset^{MO} = \mu_L^{MO} = 0$ specify an equilibrium. It is immediate that the agent's behavior is sequentially rational given these quality beliefs. Moreover, the sequences of completely mixed behavior $\beta^H(n) = (\frac{n-1}{n}, \ldots, \frac{n-1}{n})$ and $\beta^L(n) = (\frac{1}{n}, \ldots, \frac{1}{n})$ converge towards the real behavior and support the specified quality beliefs (see (1) and (2)). Since the support of the induced quality belief lottery is as in regime MT, it follows from Corollary 1 that the principal is indifferent between the two lotteries. Hence, the best quality belief lottery that can be induced in regime MO cannot be worse than that induced in MT.

It remains to show that no better quality belief lottery can be induced in regime MO. We distinguish three cases: (i) Suppose $\exists s', s'' \in \mathcal{S} : \beta_{s'}^L < 1$ and $\beta_{s''}^H < 1$. This can only be consistent with sequential rationality if neither μ_L^{MO} nor μ_H^{MO} is strictly higher than $\mu_\emptyset^{MO}$. Since the agent can ensure himself quality belief $\mu_\emptyset^{MO}$, quality belief $\mu_\emptyset^{MO}$ realizes with probability 1. By Bayes' Law, $\mu_\emptyset^{MO} = \mu_0$. By Corollary 1, this is weakly worse for the principal than the quality belief distribution induced in regime MT. (ii) Suppose $\beta^L = \mathbf{1}$. By (1), $\mu_L^{MO} = 0$. Because of mandatory participation, test result H is obtained with probability $(\mu_0 + (1 - \mu_0)\varepsilon)$. Sequential rationality requires that a single quality belief realizes when the test result is H, say μ'. By Bayes' Law, $0 + (\mu_0 + (1-\mu_0)\varepsilon)\mu' = \mu_0$ such that $\mu' = \mu_H^{MT}$. Hence, the same quality belief distribution is induced as in regime MT. (iii) Suppose $\beta^H = \mathbf{1}$. By (1), $\mu_H^{MO} = \mu_H(1, 1) = \mu_H^{MT}$. Because of mandatory participation, test result L is obtained with probability $(1 - \mu_0)(1 - \varepsilon)$. Sequential rationality requires that a single quality belief realizes when the test result is

L, say μ''. By Bayes' Law, $(1-\mu_0)(1-\varepsilon)\mu''+(\mu_0+(1-\mu_0)\varepsilon)\mu_H^{MT} = \mu_0$. Since $(\mu_0+(1-\mu_0)\varepsilon)\mu_H^{MT} = \mu_0$, $\mu'' = 0$. Hence, the same quality belief lottery as in regime MT is induced.

Consider regime VO. Participation behavior $\alpha = \mathbf{1}$ and revelation behavior $\beta^H = \mathbf{1}$, $\beta^L = \mathbf{0}$ together with quality beliefs $\mu_H^{VO} = \mu_H^{MT}$ and $\mu_\emptyset^{VO} = \mu_L^{VO} = 0$ specify an equilibrium. It is immediate that the agent's behavior is sequentially rational given these quality beliefs. Moreover, the sequences of completely mixed behavior $\alpha(n) = (\frac{n-1}{n},\ldots,\frac{n-1}{n})$, $\beta^H(n) = (\frac{n-1}{n},\ldots,\frac{n-1}{n})$ and $\beta^L(n) = (\frac{1}{n},\ldots,\frac{1}{n})$ converge towards the real behavior and support the specified quality beliefs (see (1) and (2)). Since the support of the induced quality belief lottery is as in regime MT, it follows from Corollary 1 that the principal is indifferent between the two lotteries. Hence, the best quality belief lottery that can be induced in regime VO cannot be worse than that induced in MT.

It remains to show that no better quality belief lottery can be induced in regime VO. If $\alpha = \mathbf{1}$, then the above argument applies. Suppose $\exists s \in \mathcal{S} : \alpha_s < 1$ in equilibrium. Let $s' = \max\{s \in \mathcal{S} | \alpha_s < 1\}$. We distinguish two cases: (i) Suppose $\varepsilon > 0$ or $\theta_{s'} > 0$. Then the agent obtains test results H with positive probability when his private signal is s' and he participates. By sequential rationality, $\alpha_{s'} < 1$ requires $\mu_H^{VO} \leq \mu_\emptyset^{VO}$. It follows that test result H might only be revealed if $\mu_H^{VO} = \mu_\emptyset^{VO}$. Moreover, if test result L was revealed with positive probability, Bayes' Law would imply $\mu_L^{VO} = 0$. It follows from sequential rationality that test result L is revealed with probability zero. Hence, quality belief $\mu_\emptyset^{VO}$ realizes with probability one. By Bayes' Law, $\mu_\emptyset^{VO} = \mu_0$. By Corollary 1, this is weakly worse for the principal than the quality belief distribution induced in regime MT. (ii) Suppose $\varepsilon = 0$ and $\theta_{s'} = 0$. Then the agent learns test result L with positive probability and test result H with probability zero. If test result L was revealed with positive probability, Bayes' Law would imply $\mu_L^{VO} = 0$. It follows from sequential rationality that also test result L is revealed with probability zero. Hence, quality belief $\mu_\emptyset^{VO}$ realizes with probability one. By Bayes' Law, $\mu_\emptyset^{VO} = \mu_0$. By Corollary 1, this is weakly worse for the principal than the quality belief distribution induced in regime MT. q.e.d.

Proof of Lemma 2

(a) The case of an uninformative private signal is captured by $S = 1$. It follows from (1) and (2) that for any sequence of strictly mixed behavior $\alpha(n)$, $\lim_{n\to\infty} \mu_\emptyset(\alpha(n), \mathbf{1}, \mathbf{1}) = \theta_1$, $\lim_{n\to\infty} \mu_L(\alpha(n), \mathbf{1}) = 0$ and $\lim_{n\to\infty} \mu_H(\alpha(n), \mathbf{1}) = \frac{\theta_1}{\theta_1+(1-\theta_1)\varepsilon}$. Hence, in any equilibrium $\mu_\emptyset^{VT} = \theta_1$, $\mu_L^{VT} = 0$ and $\mu_H^{VT} = \frac{\theta_1}{\theta_1+(1-\theta_1)\varepsilon}$. If the agent does not participate, the principal correctly infers his ad interim probability of being good, θ_1. If he participates, he incurs a fair quality belief lottery, i.e. a quality belief lottery with expected quality belief θ_1 but positive variance. It follows from strict concavity of u that the agent has a strict incentive not to participate such that the induced quality belief lottery is degenerate.

(b) The case of a perfect private signal is captured by $S = 2$, $\theta_1 = 0$ and $\theta_2 = 1$. It is easily checked that participation behavior $\alpha = \mathbf{1}$ together with quality beliefs $\mu_H^{VT} = \frac{p_2}{p_2+p_1\varepsilon}$, $\mu_L^{VT} = 0$ and $\mu_\emptyset^{VT} = 0$ specify an equilibrium in regime VT and that the induced quality belief distribution is as in regime MT. It remains to show that no better quality belief distribution can be induced. If $\alpha \in \{\mathbf{0}, \mathbf{1}\}$ or $\varepsilon = 0$, it is obvious that no better quality belief distribution can be induced. Suppose $\alpha \neq \{\mathbf{0}, \mathbf{1}\}$ and $\varepsilon > 0$. We distinguish four cases: (i) Suppose $\alpha_1 \in (0, 1)$. For any sequence of strictly mixed behavior $\alpha(n)$ converging towards α, we obtain from (1) $\lim_{n\to\infty} \mu_H(\alpha(n), \mathbf{1}) = \frac{p_2\alpha_2}{p_2\alpha_2+p_1\alpha_1\varepsilon} = \mu_H^{VT}$ and

$\lim_{n\to\infty}\mu_L(\alpha(n),\mathbf{1}) = 0 = \mu_L^{VT}$. Sequential rationality requires $u(\mu_\emptyset^{VT}) = \varepsilon u(\mu_H^{VT}) + (1-\varepsilon)u(\mu_L^{VT})$. It follows from this condition and $\mu_H^{VT} > \mu_L^{VT}$ that the agent has a strict incentive to take the test when $s = 2$. Hence, $\alpha_2 = 1$. For $\alpha_2 = 1$ we obtain from (2) that for any sequence of strictly mixed behavior $\alpha(n)$ converging towards α, $\lim_{n\to\infty}\mu_\emptyset(\alpha(n),\mathbf{1},\mathbf{1}) = 0$ contradicting that the agent is indifferent between taking and not taking the test when $s = 1$. Hence, there exists no equilibrium with $\alpha_1 \in (0,1)$. (ii) Suppose $\alpha_2 \in (0,1)$. Sequential rationality requires that $\mu_\emptyset^{VT} = \mu_H^{VT}$. If $\alpha_1 > 0$, we obtain from (1) that for any sequence of strictly mixed behavior $\alpha(n)$ converging towards α, $\lim_{n\to\infty}\mu_L(\alpha(n),\mathbf{1}) = 0$. It follows that an agent with $s = 1$ has a strict incentive not to take the test, contradicting $\alpha_1 > 0$. Hence, $\alpha_1 = 0$. But when $\alpha_1 = 0$, we obtain from (1) and (2) that for any sequence of strictly mixed behavior $\alpha(n)$ converging towards α, $\lim_{n\to\infty}\mu_H(\alpha(n),\mathbf{1}) = 1 = \mu_H^{VT}$ and $\lim_{n\to\infty}\mu_\emptyset(\alpha(n),\mathbf{1},\mathbf{1}) = \frac{p_2(1-\alpha_2)}{p_2(1-\alpha_2)+p_1} = \mu_\emptyset^{VT} < 1$ contradicting $\mu_\emptyset^{VT} = \mu_H^{VT}$. Hence, there exists no equilibrium with $\alpha_2 \in (0,1)$. (iii) Suppose $\alpha_1 = 0$ and $\alpha_2 = 1$. Then we obtain from (1) and (2) that for any sequence of strictly mixed behavior $\alpha(n)$ converging towards α, $\lim_{n\to\infty}\mu_H(\alpha(n),\mathbf{1}) = 1$ and $\lim_{n\to\infty}\mu_\emptyset(\alpha(n),\mathbf{1},\mathbf{1}) = 0$. Since an agent with $s = 1$ passes the test with positive probability, this contradicts optimality of $\alpha_1 = 0$. (iv) Suppose $\alpha_1 = 1$ and $\alpha_2 = 0$. Then we obtain from (1) and (2) that for any sequence of strictly mixed behavior $\alpha(n)$ converging towards α, $\lim_{n\to\infty}\mu_H(\alpha(n),\mathbf{1}) = \lim_{n\to\infty}\mu_L(\alpha(n),\mathbf{1}) = 0$ and $\lim_{n\to\infty}\mu_\emptyset(\alpha(n),\mathbf{1},\mathbf{1}) = 1$. This contradicts optimality of $\alpha_1 = 1$. q.e.d.

Proof of Lemma 3

Suppose an equilibrium with separation of private information exists, i.e., $\alpha_{s'} \neq \alpha_{s''}$ for some s', s''. Taking as given the principal's quality beliefs, the expected utility for an agent with private signal s when taking the test is $(\theta_s + (1-\theta_s)\varepsilon)u(\mu_H)$. The utility from not taking the test is $u(\mu_\emptyset)$. In any equilibrium in which $\alpha_{s'} \neq \alpha_{s''}$ for some s', s'', there is an $\tilde{s}$ such that $\alpha_s = 1$ for all $s > \tilde{s}$ and $\alpha_s = 0$ for all $s < \tilde{s}$, as the expected utility from taking the test is strictly increasing in s. Hence, μ_H is bounded from below by the expectation of θ_s conditional on $\theta_s \geq \theta_{\tilde{s}}$ and $\mu_\emptyset$ is bounded from above by the expectation of θ_s conditional on $\theta_s \leq \theta_{\tilde{s}}$. Hence, there is a discrete difference between μ_H and $\mu_\emptyset$ in any equilibrium with separation of private information. As ε goes to one, the expected utility from taking the test goes to $u(\mu_H)$ for all s. To deter participation for some s in the case that ε is high, μ_H has to go to $\mu_\emptyset$ which is not possible in an equilibrium with separation of private information. Hence, no such equilibrium exists if ε is too high. q.e.d.

Proof of Lemma 4

Consider a test with accuracy ε in regime VT. In any equilibrium with $\alpha \neq \mathbf{0}$, $\mu_L = 0$, $\mu_\emptyset$ is bounded from below by θ_1, and μ_H is bounded from above by 1. A necessary condition for the agent to participate in the case $s = 1$ is $(\theta_1 + (1-\theta_1)\varepsilon u(1)) \geq u(\theta_1)$. As the agent is information-averse, there exists a $\tilde{\varepsilon} > 0$ such that the former condition is violated if $\varepsilon < \tilde{\varepsilon}$.

As ε goes to zero, both μ_H^{VT} and μ_H^{MT} go to one. At least when $s = 1$, the agent does not take the test. As a consequence, the intermediate quality belief $0 < \mu_\emptyset \leq \mu_0 < 1$ realizes with strictly positive probability. It follows from Corollary 1 that if $\underline{\varepsilon}$ is close to zero, when choosing the most accurate test, the principal is strictly better off in regime MT than in regime VT. She can avoid the positive probability mass on $\mu_\emptyset$ only by inducing full participation, which can be achieved by increasing ε. However, the induced quality belief lottery can also be induced in regime MT by choosing the same test accuracy, and is strictly worse than the lottery induced by the most accurate test in regime MT. q.e.d.

References

[1] Faure-Grimaud, Antoine, Eloic Peyrach and Lucia Quesada, The ownership of ratings, *RAND Journal of Economics* 40(2), 234-257.

[2] Farhi, Emmanuel, Josh Lerner and Jean Tirole (2010): "Fear of Rejection? Tiered Certification and Transparency", working paper.

[3] Lizzeri, Alessandro (1999): Information revelation and certification intermediaries, *RAND Journal of Economics* 30(2), 214-231

[4] Rosar, Frank and Elisabeth Schulte (2010): "Imperfect private information and information-generating mechanisms", working paper.