Designing the Optimal Menu of Tests

Nathan Hancart*

November, 2022

CLICK HERE FOR LATEST VERSION

Abstract

A decision-maker must accept or reject a privately informed agent. The agent always wants to be accepted, while the decision-maker wants to accept only a subset of types. The decision-maker has access to a set of feasible tests and, prior to making a decision, requires the agent to choose a test from a menu, which is a subset of the feasible tests. By offering a menu, the decision-maker can use the agent's choice as an additional source of information. I characterise the decision-maker's optimal menu for arbitrary type structures and feasible tests. I then apply this characterisation to various environments. When the domain of feasible tests contains a most informative test, I characterise when only the dominant test is offered and when a dominated test is part of the optimal menu. I also characterise the optimal menu when types are multidimensional or when tests vary in their difficulty.

*Nathan Hancart: Department of Economics, University College London, 30 Gordon St, London WC1H 0AN, UK, nathan.hancart.16@ucl.ac.uk. I am grateful to Ran Spiegler and Vasiliki Skreta for their guidance and support. I thank Deniz Kattwinkel for his availability and help in this project. Finally, I would also like to thank Anna Becker, Alex Clyde, Martin Cripps, Duarte Gonçalves, Amir Habibi, Philippe Jehiel and the theory group at UCL and Penn State for their helpful comments.

1 Introduction

In many economic settings, decision-makers (DMs) rely on tests to guide their actions. Universities use standardised tests as part of their admission process, firms interview job candidates before they hire them and regulators test products prior to authorisation. In these examples, the DM is deciding whether to accept an agent and his preferences depend on some private information held by the agent: the ability of the student, the productivity of the candidate or the quality of the product. Ideally, the DM would want to use a fully revealing test, but this is often not feasible and thus the ability to learn from the test outcome alone is limited. However, there is an additional channel the DM can learn from: he can offer a menu of tests and let the agent *choose* which test to take. The DM can then use the agent's choice as an additional source of information.

Depending on the setting, constraints on the testing ability can take many different forms. For instance, a hiring firm is constrained by the amount of time and resources it can allocate to the selection process; most universities have to use externally provided tests like the SAT or the GRE for their admission procedures; and medicine regulatory agencies face both technological and ethical constraints when authorising new drugs.

In this paper, I study the DM's optimal design of a menu of tests and provide conditions under which the DM learns from both the test choice and outcome. I first develop general tools for characterising the optimal menu of tests for arbitrary domains of feasible tests. I then apply these tools to natural economic applications and determine which tests are part of the optimal menu and how it depends on their properties. I provide conditions under which it is optimal to include a strictly less informative test in the optimal menu and when it is not. I also characterise the optimal menu when tests vary in their difficulty and when each test can identify only one dimension of the agent's private information.

I study a DM who has to accept or reject a privately informed agent. While the DM wants to accept a subset of types (the A-types) and reject the others (the R-types), the agent always wants to be accepted. The domain of feasible tests is an exogenously given set of Blackwell experiments. The DM designs a menu of tests, a subset of the feasible tests, from which the agent chooses one. For example, a university can let a student decide whether he takes a standardised test in its admission procedure. A regulator can let a pharmaceutical company design the clinical trials when authorising a new drug. After observing, the choice of test and its result, the DM decides whether or not to accept.

I characterise optimal menus for arbitrary type structures and domain of feasible tests. Specifically, I show that the optimal menu and strategies are the outcome of an auxiliary max-min problem. This result greatly simplifies the analysis: rather than comparing equilibria under different menus to determine the optimal one, it is enough to focus on a simpler optimisation problem. In that max-min problem, the accept types and the DM maximise the DM's utility by choosing a test and a decision rule, while the reject types minimise it by selecting an accept type to mimic. Tests chosen in that problem correspond to the optimal menu. Moreover, I show there is no value to commitment: regardless of whether the DM is able commit ex-ante to an acceptance rule, the optimal menu and strategies remain the same. Finally, I find that it is without loss to consider menus with as many tests as accept types. This implies that if there is only one type the DM wants to accept, it is sufficient to consider menus with a single test and the only information used is the test's outcome.

In Section 4, I use Theorem 1 to determine which tests are part of the optimal menu in three natural economic applications. In Section 4.1, I consider a domain of feasible tests containing a dominant one, in the sense of Blackwell's (1953) informativeness order. One example of this environment is a university considering whether to allow students to opt out of a standardised test like the SAT when applying. This is effectively offering a menu with the SAT and an uninformative test.

In Lemma 1, I show that the most informative test is always part of an optimal menu. I then provide conditions under which a dominated test is part of the optimal menu when tests are binary, i.e., pass-fail tests. In this case, types can be ordered by how likely they are to generate the pass signal in the most informative test. I show that the optimal menu always includes a strictly less informative test if, and only if, the DM's payoff is enclosed. This corresponds to the DM wanting to accept at least the worst and the best performer on the test. On the other hand, for any prior an optimal menu contains only the most informative test if, and only if, the DM's payoff is single-peaked with respect to that order. This corresponds to the DM willing to accept either only high types, only low types or only intermediate types, as measured by their performance on the test. Failure of single-peakness can occur for example when the most informative test does not test all relevant dimensions or only tests a proxy of the relevant dimension.

In the case where the domain of feasible tests contains a dominant tests and they generate more than two signals, the results extend as follows. If there exists a subset of signals where single-peakness is violated, there exists a less informative test that is part of the optimal menu for some prior. On the other hand, if the environment is one-dimensional, in the sense that all the tests satisfy the monotone likelihood ratio property and the DM wants to accept any type above a threshold, only the most informative test is offered.

In the first application, I considered a domain of feasible tests where tests can be ordered by their informativeness. In Section 4.2, I consider one-dimensional environments where feasible tests are ordered by their difficulty. For example, the DM could be a regulator deciding how hard a compliance test is before authorising a product. The testing technology is a set of pass-fail tests and varying the difficulty of a test changes which types it identifies better. A more difficult test is informative when it is passed, as only high types are likely to produce a high grade but it is less informative when it is failed. In this case, I show again that a singleton menu is optimal.

In previous sections, I show that for natural specifications of one-dimensional environments, a singleton menu is optimal. I then turn to multidimensional environments. For example, a hiring firm could be guided by two considerations, the candidate's technical and managerial skills and specialise the interview on either dimension. More generally, I assume that the agent's type has two components and each test is informative about only one of them.[1] Offering tests for both dimensions allows A-types that perform badly in one dimension to select the test where they perform best. I show that the optimal menu contains both tests for any prior if and only if the DM wants to accept any type that performs well in at least one dimension. This would be the case if the hiring firm would want to hire a candidate with high technical skills but no managerial skills and vice-versa. On the other hand, if the firm cares about both dimension simultaneously, then, for some priors, it uses only one test.

In Section 5, I move beyond specific applications and give a general condition on the DM's preferences and tests available that guarantees that a test is part of an optimal menu. I also show the necessary and sufficient condition on tests for the DM to always make the right decision in the optimal menu.

Finally, in Section 6, I show that the model can be easily extended to allow for communication. I model communication as an additional cheap-talk message on top of the test choice. For example, it could be a cover letter where the candidate can freely communicate with the DM when applying for a job or to university. A characterisation as in Theorem 1 also holds. I also show that in this case, it is irrelevant for the outcome of the game who chooses the test, the DM or the agent.

[1] The results extend easily to more than two dimensions.

Relation to the literature

This paper relates to both the literature on strategic disclosure and mechanism design with evidence and the literature on information design without commitment. The strategic disclosure literature studies information provision by privately informed players. In these papers, information provision is usually modelled with hard evidence (e.g., Grossman, 1981; Milgrom, 1981; Dye, 1985; Milgrom, 2008). Hard evidence is a particular kind of test that takes a deterministic form: the agent can provide evidence that he belongs to a certain subset of types. Another difference with modelling information with evidence is that, in my language, not all types can participate in all tests. Instead, I allow arbitrary stochastic tests and all types can participate in any test. I discuss the relation between these two modelling approach in more details in Section 2.2.

Formally, my model is most closely related to Glazer and Rubinstein (2006). They also study a problem where an agent wants to persuade a DM to accept him but in their model, the agent can only present deterministic evidence about his type. They characterise the optimal mechanism that maps evidence to a decision and show that the outcome can be implemented without commitment (see also Hart et al., 2017, for similar results with other payoff structures; Sher, 2011). They also show that with deterministic evidence, the optimal decision rule is deterministic. I extend their analysis in two ways. First, Theorem 1 generalises their result on commitment to arbitrary testing technology and my characterisation result also applies in their setup. I also show that that the optimal decision rule is no longer deterministic when tests are stochastic. Second, I use the characterisation to prove general results on which test is included in the optimal menu depending on the properties of the feasible tests.

Glazer and Rubinstein (2004) study a related problem where the agent communicates with the DM who, based on the communication, verifies one dimension of a multidimensional type. In In Section 6, I extend the characterisation of Theorem 1 to allow for communication from the agent and show how to generalise some of Glazer and Rubinstein (2004)'s results on the value of commitment.

More generally, this paper relates to the mechanism design with evidence literature (e.g., Green and Laffont, 1986; Bull and Watson, 2007; Deneckere and Severinov, 2008; Koessler and Perez-Richet, 2019; Forges and Koessler, 2005; Kartik and Tercieux, 2012; Strausz, 2017). Assuming commitment from the DM, Theorem 1 characterises the optimal mechanism that maps a test choice and test outcome to an acceptance probability. Unlike most of that literature, I allow for arbitrary domain of feasible tests that include non-deterministic tests.[2] The payoff structure assumed in this paper is commonly used in this literature, e.g., in Glazer and Rubinstein (2004); Glazer and Rubinstein (2006) and special cases of Ben-Porath et al. (2019); Ben-Porath et al. (2021). The characterisation of Theorem 1 can be applied in these settings as well and thus provides a useful tool beyond the model and applications considered here.

An important focus of the literature on strategic disclosure is finding conditions under which all information is revealed in equilibrium, see e.g., Grossman (1981), Milgrom (1981), Lipman and Seppi (1995), Giovannoni and Seidmann (2007), Hagenbach et al. (2014) or Carroll and Egorov (2019). In my model, if full information is possible, it is optimal, but I also characterise the optimal choice of test when full information is not attainable. In Proposition 9, I provide the necessary and sufficient conditions for full payoff-relevant information revelation.

The other branch of literature my paper relates to is information design without sender commitment. In these papers, the agent and the DM correspond to the sender and the receiver. In particular, this paper is closer to models characterising receiver-optimal tests where the sender can choose which test to take. Rosar (2017) and Harbaugh and Rasmusen (2018) the

[2]For an example of mechanism design paper with non-deterministic tests, see Ball and Kattwinkel (2022), Ben-Porath et al. (2021).

receiver designs a test where a privately informed agent can either take the test, possibly at a cost, or take an uninformative test. In these papers, the receiver flexibly designs a test *given* that the sender has a choice. In my paper, the receiver designs the choice, i.e., the menu, given the restrictions on the feasible tests.

Other papers consider the receiver-optimal design of tests where the sender's action is partially observed or unobserved, e.g., DeMarzo et al. (2019), Deb and Stewart (2018), Perez-Richet and Skreta (2022) or Ball (2021) (note that Perez-Richet and Skreta (2022) also consider observable action). The design of the optimal test also has to take into account the strategy of the sender, however unobservable actions fundamentally changes the sender's incentives and thus how information is revealed. I discuss in Section 2.2 which results would still apply if the outcome of the tests depends on the agent's unobserved effort.[3]

Finally, this paper is related to Ely et al. (2021). They study the optimal allocation of tests from a restricted set to agents with observable characteristics. My paper can be interpreted as a problem of optimal allocation of tests with asymmetric information, thus the allocation must also respect incentive constraints.

2 Model

There is a decision-maker (DM) and an agent. The agent has a type $\theta \in \Theta$, $|\Theta| < \infty$, with a common prior $\mu \in \Delta(\Theta)$. The set of types is partitioned in two: $\Theta = A \cup R$, $A \cap R = \emptyset$. The type is private information of the agent. The DM must take an action $a \in \{0, 1\}$, accept or reject. The utilities of the DM and the agent are $v(a, \theta) = a\big(\mathbb{1}[\theta \in A] - \mathbb{1}[\theta \in R]\big)$

[3]There are also papers studying sender-optimal tests when the sender cannot fully commit to reporting the test correctly, e.g., Nguyen and Tan (2021), Lipnowski et al. (2022) or Koessler and Skreta (2022). In Boleslavsky and Kim (2018) and Perez-Richet et al. (2020), the sender can commit but there is a third agent whose effort determines respectively the state of the world or the Blackwell experiment actually performed.

and $u(a,\theta) = a$. That is, the DM wants to accept agents in A and reject agents in R. The agent always wants to be accepted. The analysis is virtually unchanged by allowing for DM's utility functions of the form $v(a,\theta) = a\nu(\theta)$ for some $\nu : \Theta \to \mathbb{R}$.

There is a finite exogenous set of test $T \subseteq \Pi \equiv \{\pi : \Theta \to \Delta X\}$, where X is some finite signal space. The conditional probabilities of test t are $\pi_t(\cdot|\theta)$. The set T captures the restriction on the DM's testing capacity. He can only perform one test from that set. A menu of test is a subset of the feasible tests, $\mathcal{M} \subseteq T$.

The timing of the game is as follows. For a menu $\mathcal{M} \subseteq T$,

1. The agent learns his type θ.
2. The agent chooses a test from the menu, denoted by $\sigma : \Theta \to \Delta\mathcal{M}$.
3. A signal x is drawn according to $\pi_t(\cdot|\theta)$.
4. The DM chooses an action based on the realised test choice and outcome, the acceptance probability denoted by $\alpha : \mathcal{M} \times X \to [0,1]$.

Beliefs of the DM are $\tilde{\mu} : \mathcal{M} \times X \to \Delta\Theta$, a probability distribution over types given an observed test and signal realisation.

The solution concept is DM-preferred weak Perfect Bayesian Equilibrium.

I write $(\alpha,\sigma) \in \text{wPBE}(\mathcal{M})$ if there is a belief $\tilde{\mu}$ where $(\alpha,\sigma,\tilde{\mu})$ is a weak PBE when the menu is $\mathcal{M}$.

The optimal design of menu solves

$$V = \max_{\mathcal{M}\subseteq T} \max_{\sigma,\alpha} \sum_{\theta\in A} \mu(\theta) \sum_{t\in\mathcal{M}} \sigma(t|\theta) \sum_{x} \alpha(t,x)\pi_t(x|\theta) - \sum_{\theta\in R} \mu(\theta) \sum_{t\in\mathcal{M}} \sigma(t|\theta) \sum_{x} \alpha(t,x)\pi_t(x|\theta)$$
$$\text{s.t. } (\alpha,\sigma) \in \text{wPBE}(\mathcal{M})$$

The inner maximisation problem selects, for a fixed menu, the DM and agent strategy to maximise the DM's payoff for a fixed menu, under the constraint that they are equilibrium strategies. The outer maximisation problem selects the best possible menu for the DM.

Notation: For any α, denote the probability of type θ to be accepted in test t by $p_t(\alpha;\theta) \equiv \sum_x \alpha(t,x)\pi_t(x|\theta)$.

Off-path beliefs: The results would exactly the same if I would take DM-preferred Sequential Equilibrium (Kreps and Wilson, 1982) as my solution concept. I comment on this in more detail in the discussion of Theorem 1.

Test restriction: The exogenous set of tests T can capture different constraints on DM's testing capacity. It could be a purely technological constraint, e.g., when choosing amongst standardised test, universities can only choose from an exogenously given set of tests (SAT, ACT, GRE, etc.). The constraint can also be on some properties of the tests that can be used, e.g., $T \subset \{\pi : \pi \text{ has the MLRP}\}$. Finally, it could come from a capacity constraint in the information processing/acquisition abilities of the DM, e.g., a limited number of sample sizes a researcher can collect or there could be a cost function associated with each experiment $C : \Pi \to \mathbb{R}$ and a maximum cost the DM can pay $c \in \mathbb{R}$, $T \subset \{\pi : c \geq C(\pi)\}$.

2.1 Example: Opting out of an admission test

Suppose a university uses some standardised test for university admission and that there are three types of students: $A = \{A1, A2\}$ and $R = \{R1\}$. Consider the testing set $T = \{t, \emptyset\}$

where $\emptyset$ is an uninformative test. The test t is described by $X = \{x_0, x_1\}$ and

$$\pi_t(x|A1) = \begin{cases} 1/2 & \text{if } x = x_0 \\ 1/2 & \text{if } x = x_1 \end{cases} \qquad \pi_t(x|R1) = \begin{cases} 1/3 & \text{if } x = x_0 \\ 2/3 & \text{if } x = x_1 \end{cases}$$

$$\pi_t(x|A2) = \begin{cases} 0 & \text{if } x = x_0 \\ 1 & \text{if } x = x_1 \end{cases}$$

Furthermore, suppose that $\mu(A1) < \frac{2}{3}\mu(R1) < \mu(A2)$.

This example can be interpreted as follows. The test t is a standardised test a university uses to get information about students, like the SAT or GRE. The signal x_1 represents a high grade and x_0 a low grade. A common concern about these tests is that they can be too easily gamed or fail to identify good students in some categories of the population (see e.g., Hubler, 2020). The parametrisation of the test t captures this phenomenon. While A_2 and R_1 are naturally ordered, in the sense that A_2 is more likely to have a good grade than R_1, A_1 corresponds to a type of student that the university wants to accept but generates a lower grade than R_1. Adding $\emptyset$ to the menu allows the student to opt out from the standardised test.

When only t is offered: The information structure and prior deliver the following best response when only t is offered,

$$\alpha(x, t) = \begin{cases} 0 & \text{if } x = x_0 \\ 1 & \text{if } x = x_1 \end{cases}$$

The acceptance probabilities of each types are then

$$p_t(\alpha; R1) = 2/3 \qquad p_t(\alpha; A1) = 1/2 \qquad p_t(\alpha; A2) = 1$$

When both t and $\emptyset$ are offered: Consider the equilibrium with the following strategies of the agent:

$$\sigma(\emptyset|R1) = \frac{\mu(A1)}{\mu(R1)} \qquad \sigma(\emptyset|A1) = 1 \qquad \sigma(t|A2) = 1$$

The student $R1$ mixes between the two tests, t and $\emptyset$, whereas $A1$ chooses $\emptyset$ with probability one and $A2$ chooses t with probability one. Note that if all types play a pure strategy, it is not possible to maintain an equilibrium where both tests are chosen. If it is the case, there is a test that is only chosen by an A-type and in equilibrium the DM must accept with probability one after any signal in that test. Thus R_1 mixes in equilibrium to make the menu $\{t, \emptyset\}$ credible.

Given the agent's strategy, the DM's strategy after t remains the same as before. When the DM observes $\emptyset$, he is indifferent between accepting and rejecting. He then mixes in a way that makes $R1$ indifferent between $\emptyset$ and t: $\alpha(x, \emptyset) = 2/3$. The resulting acceptance probabilities are

$$\mathbb{E}[p(\alpha; R1)] = 2/3 \qquad p_{\emptyset}(\alpha; A1) = 2/3 \qquad p_t(\alpha; A2) = 1$$

Types $R1$ and $A2$ have the same acceptance probabilities as before but $A1$ is accepted with strictly higher probability. Therefore, allowing to opt out strictly increases the DM's payoffs.

2.2 Discussion

Effort: The outcome of the test is independent of the agent's action. The model would go unchanged if effort is costless and observable as it could be deterred with off-path beliefs. If the effort is costless but unobservable the results would generally change. However, if signals are ordered and the DM uses a cutoff strategy, as in many natural applications, a reasonable assumption on effort would be that the higher the effort, the likelier a high signal. In this case, the agent would always have an incentive to provide high effort. See Deb and Stewart (2018) and Ball and Kattwinkel (2022) for models that takes into account both asymmetric information and moral hazard in a model of testing.

Relation to models with evidence: The model can be interpreted as a generalisation of models with evidence. The idea of these models is that each type is endowed with a set of messages that only a subset of types can send. Formally, an evidence structure is a correspondence $E : \Theta \rightrightarrows M$ for some finite set of messages M. Thus type θ can only send messages in $E(\theta)$. We can capture these models in the following way. The set of feasible test has $X = \{x_1, x_0\}$ and for all $m \in M$, $\pi_m(x_1|\theta) = 1 \Leftrightarrow \theta \in E^{-1}(m)$. Thus a test m perfectly reveals whether θ is in $E^{-1}(m)$ or in $\Theta \setminus E^{-1}(m)$. In a model with evidence, a type θ can never reveal he is in $\Theta \setminus E^{-1}(m)$ for a message $m \notin E(\theta)$. However, in the testing model, we can always incentivise any type to not choose such a test by setting $\alpha(x_0, m) = 0$ for all m. This strategy could be justified because (x_0, m) would always be off-path. Alternatively, we can set this restriction on α directly and Theorem 1 would still hold.

3 Characterisation of the optimal menu

In this section, I show that the value of the optimal menu is characterised by a solution of a max-min problem.

Let $s : A \to \Delta T$ and $m : R \to \Delta A$ and abusing notation, let $\alpha : T \times X \to [0, 1]$ and

$$\begin{aligned} v(\alpha, s, m) &\equiv \sum_{\theta \in A} \sum_{t \in T} s(t|\theta) \Big[\mu(\theta) p_t(\alpha; \theta) - \sum_{\theta' \in R} \mu(\theta') m(\theta|\theta') p_t(\alpha; \theta') \Big] \quad (1) \\ &= \sum_{\theta \in A} \sum_{t \in T} \mu(\theta) s(t|\theta) p_t(\alpha; \theta) - \sum_{\theta' \in R} \mu(\theta') \sum_{\theta \in A} m(\theta|\theta') \sum_{t \in T} s(t|\theta) p_t(\alpha; \theta') \end{aligned}$$

The function s can be interpreted as A-types choosing a test, m as R-types choosing an A-type to mimic, α as the DM accepting the agent after a test and signal realisation. The function v is then the DM's expected payoffs from a distribution over tests induced by the pair (s, m). I explain these objects in more detail in the discussion of Theorem 1.

Theorem 1. *The value of an optimal menu is*

$$V = \max_{\alpha, s} \min_{m} v(\alpha, s, m)$$

There is $(\alpha, s) \in \arg\max_{\tilde{\alpha}, \tilde{s}} \min_{\tilde{m}} v(\tilde{\alpha}, \tilde{s}, m)$ *and* $m \in \arg\min_{\tilde{m}} \max_{\tilde{\alpha}} v(\tilde{\alpha}, s, \tilde{m})$ *with* $s(\cdot|\theta)$ *in pure strategy for all* $\theta \in A$ *such that*

- *an optimal menu is* $\mathcal{M} = \cup_{\theta \in A} \operatorname{supp} s(\cdot|\theta)$*,*
- *and the strategies in this optimal menu are*
  - *for* $\theta \in A : \sigma(t|\theta) = s(t|\theta)$
  - *for* $\theta' \in R : \sigma(t|\theta') = \sum_{\theta \in A} m(\theta|\theta') s(t|\theta)$
  - *the DM's strategy is* α*.*

Moreover, the DM does not benefit from committing to α*.*

All proofs are relegated to the appendix.

Theorem 1 provides a characterisation of the optimal menu in terms of an auxiliary max-min problem. The fact that an optimal menu is a solution to an optimisation problem gives us a powerful tool to test equilibria. Indeed, it is not necessary to compare equilibria across menus to establish that a menu is not optimal. It is enough to find that $(\tilde{\alpha}, \tilde{s})$ such that

$$\min_{m} v(\alpha, s, m) < \min_{m} v(\tilde{\alpha}, \tilde{s}, m)$$

to show that (α, s, m) does not constitute an optimal menu without having to worry whether $(\tilde{\alpha}, \tilde{s})$ is optimal.

To understand the structure of this max-min problem better, consider the objective function v for a fixed α. This can be interpreted as a zero-sum game where where the maximiser, the A-types, chooses $s : A \to \Delta T$ and the minimiser, the R-types, chooses $m : R \to \Delta A$. Consider the payoffs of a given A-type θ choosing test t and a given R-type, θ', choosing an A-type $\tilde{\theta}$:

$$\text{for } \theta \in A \text{ choosing } t,\ \mu(\theta)p_t(\alpha;\theta) - \sum_{\theta' \in R} \mu(\theta')m(\theta|\theta')p_t(\alpha;\theta')$$

$$\text{for } \theta' \in R \text{ choosing } \tilde{\theta},\ \mu(\theta') \sum_t s(t|\tilde{\theta})p_t(\alpha;\theta') - \sum_{\theta \in A, t} \mu(\theta)s(t|\theta)p_t(\alpha;\theta)$$

Note that in the payoffs of the R-type, his strategy, the choice of $\tilde{\theta}$, only affects the first part of the payoffs. So the R-type is effectively trying to maximise his probability of being accepted.

On the other hand, the A-type maximise a modified version of their utility where they maximise their probability of being accepted while being penalised every time a R-type mimics them and is accepted. The A-types' utility is thus modified to align it with the DM's payoffs.

The strategies of that game induce a distribution over tests for each type. The A-types get the distribution over test they choose and the R-types the distribution of the A-types they choose to mimic. Theorem 1 shows that the tests chosen by the A-types are actually the tests chosen in the optimal menu. Moreover, the A-types play a pure strategy.

The outcome of the max-min problem delivers an incentive compatible outcome for the following reason. Any A-type when choosing a test in the max-min problem is trading-off his acceptance probability and the acceptance probability of the R-types mimicking him. Now note that any deviation from an A-type to another A-type strategy always weakly reduces the R-types' payoffs as they have less distribution over tests they can mimic. Thus, if s is not incentive compatible, an A-type could both increase his acceptance probability and the acceptance probability of all the R-types, contradicting that $s \in \arg\max_{\tilde{s}} \min_{\tilde{m}} v(\alpha, \tilde{s}, \tilde{m})$.

Theorem 1 also shows that commitment has no value. I interpret this result as a hierarchy over sources of learning. The DM has two sources of information, the "hard information" from the test results and the endogenously created information from the choice of test. When the DM can commit to a strategy, he can "sacrifice" payoffs from the test result by not best replying, in order to create separation of types through the test choice. By showing that the DM always best replies, even when he can commit, I show that he should always prioritise the hard information over creating endogenous information through the test choice.

That commitment has no value in this game comes from the zero-sum structure of the characterisation. Because a minimax theorem holds, this implies that the order of moves do not matter in this game: the DM has the same payoffs if he moves first or last.

Note that if the solution concept is DM-preferred Sequential Equilibrium (SE) (Kreps and Wilson, 1982), Theorem 1 would also hold. If all tests have full support, then all signals are on-path and the PBE and SE coincide. If some tests do not have full support, then I can always assume that the trembling of R-types is more likely than the trembling of A-types. Then, the DM's off-path beliefs after the pair (t, x) are that the type is an A-type if the support of A- and R-types do not coincide and that the type is an R-type otherwise. This guarantees that if an A-type finds it profitable to deviate in the menu game, he would also find it profitable in the zero-sum game as no R-type would have an incentive to mimic him.

Finally, Theorem 1 gives an upper bound on the number of tests needed in an optimal menu. If A-types are playing a pure strategy and R-types only use tests A-types use, then the number of tests used is at most $|A|$.

Corollary 1. *The number of tests used in the optimal menu is at most $|A|$.*

An immediate corollary is also that if there is only one type the DM would like to accept an optimal menu is to use only one test. In particular, this results shows that in a binary state environment, the optimal mechanism uses only one test, no matter what the available set of test is.

Corollary 2. *Suppose $|A| = 1$. Then for any T, there is an optimal menu that uses only one test.*

3.1 Sketch of proof Theorem 1

To prove Theorem 1, I first introduce *mechanisms*. A (direct) mechanism is a mapping $\tilde{\sigma} : \Theta \to \Delta T$, a function from types to distribution over tests. In the optimal mechanism design problem, the DM chooses $\tilde{\sigma}$ to maximise the his payoffs. The DM only observes the realised test and signal, thus the definition of his strategy is unchanged. The agent's strategy is now to report a type into the mechanism. The solution concept is still DM-preferred PBE. Standard arguments show that without loss of generality we can restrict attention to direct truthful mechanism. The designer's problem is

$$\begin{aligned}
\tilde{V} = \max_{\tilde{\sigma},\alpha} \; & \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) - \sum_{\theta \in R} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) \\
\text{s.t. } & \sum_t (\tilde{\sigma}(t|\theta) - \tilde{\sigma}(t|\theta')) p_t(\alpha;\theta) \geq 0 \text{ for all } \theta, \theta' \\
& \alpha \in BR(\tilde{\sigma})
\end{aligned}$$

where the first constraint is the agent's incentive compatibility constraint, the second is a feasibility constraint and $\alpha \in BR(\tilde{\sigma})$ means that the strategy α is a best-response to some beliefs consistent with the mechanism. We have $\tilde{V} \geq V$, i.e., the value of the optimal mechanism is larger than the value of the optimal menu, as imposing a menu is simply restricting the class of mechanism the designer could use.

The first part of the proof shows that $\tilde{V} = \max_{\alpha,s} \min_m v(\alpha, s, m)$. The second part shows that the optimal mechanism can be implemented by posting a menu of tests.

To show that $\tilde{V} = \max_{\alpha,s} \min_m v(\alpha, s, m)$, I characterise the optimal mechanism when the DM commits to α. The designer's problem becomes

$$\begin{aligned} \tilde{V}(\alpha) = \max_{\tilde{\sigma}} & \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha; \theta) - \sum_{\theta \in R} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha; \theta) \\ \text{s.t.} & \sum_t (\tilde{\sigma}(t|\theta) - \tilde{\sigma}(t|\theta')) p_t(\alpha; \theta) \geq 0 \text{ for all } \theta, \theta' \end{aligned}$$

The notation $\tilde{V}(\alpha)$ indicates the designer's problem when the DM has committed to the strategy α. I relax the problem further and only require that the IC constraints of R-types deviating to A-types hold. Deviations in the relaxed problem correspond to mimicking strategies:

$$\sum_t \tilde{\sigma}(t|\theta') p_t(\alpha; \theta') \geq \max_{m(\cdot|\theta')} \sum_{\theta \in A} m(\theta|\theta') \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha; \theta') \text{ for all } \theta' \in R$$

In the relaxed problem, the remaining IC constraints must bind, for otherwise the DM could set $\tilde{\sigma}(t|\theta') = \max_{m(\cdot|\theta')} \sum_{\theta \in A} m(\theta|\theta') \tilde{\sigma}(t|\theta)$ and strictly benefit. Thus plugging in the binding IC constraints in the objective function, we get

$$\begin{aligned} \tilde{V}(\alpha) &= \max_{\tilde{\sigma}} \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha; \theta) - \sum_{\theta' \in R} \mu(\theta') \max_{m(\cdot|\theta')} \sum_{\theta \in A} m(\theta|\theta') \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha; \theta') \\ &= \max_s \min_m v(\alpha, s, m) \end{aligned}$$

The second part shows that the optimal mechanism can be implemented by posting a menu of tests. The way the proof proceeds is by showing that there is $(\alpha^*, s^*) \in \arg\max \min_m v(\alpha, s, m)$, where s^* is a pure strategy for all $\theta \in A$. If this is the case, then we can take the menu of tests as the support of tests in the optimal mechanism. Each type $\theta \in A$ is better off choosing "his" test as choosing another one would violate the incentive compatibility constraints. Types $\theta' \in R$ possibly have a randomised allocation but they are indifferent between any tests they are allocated to. Indeed, their randomised allocation corresponds to a mixed strategy in the auxiliary game where they are maximising their probability of being accepted, just like in the menu-game.

To understand why s^* must be a pure strategy, note that $s *^{\in} \arg\max_s \min_m v(\alpha, s, m)$. It is without loss of generality to select a minimiser in pure strategy as v is linear in m. In particular, if a given $\theta \in A$ uses a mixed strategy, this strategy must be optimal against the selection of the minimiser where all $\theta' \in R$ for whom mimicking θ is a best-response, mimic that type with probability one. Let $Z \subseteq R$ such that for all $\theta' \in R$ mimicking θ is a best-response. Thus for all $\theta' \notin Z$, mimicking another $\tilde{\theta} \in A$ is strictly optimal. If $\theta \in A$ were to slightly perturb his mixed strategy, the set of R-types mimicking him would weakly decrease. Thus the only way that this perturbation is not optimal is if the payoffs from all tests in the support given Z are the same. This is a knife-edge case, and thus for a dense subset of payoffs, pure strategies are optimal. A convergence argument shows it is the case for all payoffs.

4 Applications

4.1 Optimal menu with Blackwell dominant test

It is common in applications that the DM has access to a most informative test. This can be because the choice is simply between a test and opting out of the test like in the SAT example. It can also come from the structure of the constraints. For example, the DM could have a time budget to conduct an interview. The more time the interview takes, the more informative it is. Another possibility is that the DM can easily make a test less informative by simply not conducting part of the test. If a test is composed of a series of questions, the DM can ignore some of them.

I will use Blackwell (1953)'s notion of informativeness.

Definition 1 (Blackwell (1953))**.** *A test t is more informative than t', $t \succeq t'$, if there is function $\beta : X \times X \to [0,1]$ such that for all $x' \in X$, $\sum_x \beta(x,x')\pi_t(x|\theta) = \pi_{t'}(x'|\theta)$ for all $\theta \in \Theta$ and for all $x \in X$, $\sum_{x'} \beta(x,x') = 1$.*

I call a test t a dominant test if $t \succeq t'$ for all $t' \in T$. If a test is more informative than another then in any decision problem, i.e., a pair of utility function and a prior, using the more informative test yields higher expected utility. A first important fact we will record here is that if there is a most informative test, then it is part of an optimal menu.

Lemma 1. *If there is $t \in T$ such that $t \succeq t'$ for all $t' \in T$, then there is an optimal menu that includes t.*

This lemma follows from the max-min characterisation of Theorem 1 and the properties of dominant test. Indeed, if we find a menu where the dominant test t is not included, we can modify the DM's strategy such that one A-type is accepted with higher probability than the test he is choosing, say t', and all R-types are accepted with lower probability than in t'. Then this A-type has a profitable deviation to t.

As we have seen in the SAT example in Section 2.1, it can be optimal to add a strictly less informative in the optimal menu. I first focus on binary signals environment, $X = \{x_0, x_1\}$. Let t be the most informative test in T. When signals are binary, we can order the types by their likelihood of generating signal x_1: $\theta \geq_t \theta' \Leftrightarrow \pi_t(x_1|\theta) \geq \pi_t(x_1|\theta')$.[4] I characterise the optimal menu for different payoff function of the DM.

Definition 2. *The DM's preferences are single-peaked given the order $\geq$ on Θ if there is $\theta_1, \theta_2 \in A$ such that $A = \{\theta : \theta_1 \leq \theta \leq \theta_2\}$.*

Preferences are single-peaked if the DM only wants to either only accept high types, only low types or only intermediate types, where the order is determined by the performance of types on the test. Preferences are not single-peaked whenever it is possible to find $A_1, A_2 \in A$ and $R_1 \in R$ such that $A_1 <_t R_1 <_t A_2$. This was for example the case in the SAT example in Section 2.1.

We get the following characterisation.

Proposition 1. *Let $X = \{x_0, x_1\}$. Suppose there is $t \in T$ such that $t \succeq t'$ for all $t' \in T$ and let $\geq_t$ on Θ be the order implied by t.*

The singleton menu $\{t\}$ is optimal for any $\mu \Leftrightarrow$ the DM's preferences are single-peaked given $\geq_t$.

From Lemma 1, the most informative test is part of the optimal menu. Whenever the DM's preferences are single-peaked, if the most informative test is included in the menu, the unique resulting equilibrium is one where all types choose the most informative test. The key argument in the analysis is noting that $p_t(\alpha;\theta) - p_{t'}(\alpha;\theta)$ is single-crossing in θ with respect to

[4]Note that given that tests are binary, this is equivalent to ordering type by the likelihood ratio, $\frac{\pi(x_1|\theta)}{\pi(x_0|\theta)}$.

the order $\geq_t$, for any α. When preferences are single-peaked, we can use the single-crossing condition and properties of tests satisfying the monotone likelihood ratio property to show that there is a unique equilibrium where only t is chosen.

On the other hand, if the preferences are not single-peaked, there is a prior where offering even a completely uninformative test with the most informative test is strictly better for the DM. To illustrate, consider three types $A_1, A_2 \in A$ and $R_1 \in R$ such that $A_1 <_t R_1 <_t A_2$. Suppose the prior is such that if only t is offered, the DM accepts after x_1 and rejects after x_0. The DM can then offer an uninformative test where the probability of being accepted makes R_1 indifferent but is strictly preferred by A_1. This constitutes a deviation in the max-min problem. This reasoning can be used to show that including a less informative test is always beneficial whenever the DM's payoff is *enclosed*: there is $\theta_1, \theta_2 \in A$ such that $\theta_1 <_t \theta <_t \theta_2$ for any $\theta \neq \theta_1, \theta_2$.

Proposition 2. *Let $X = \{x_0, x_1\}$. Suppose there is $t \in T$ such that $t \succeq t'$ for all $t' \in T$ and let $\geq_t$ on Θ be the order implied by t.*

If the DM's preferences are enclosed given $\geq_t$ $\Leftrightarrow$ the DM's payoffs are higher in the menu $\{t, t'\}$ than in $\{t\}$ for any μ and $t' \in T$.

The ideas of Proposition 1 and Proposition 2 can be partially extended to more than two signals. First, if all tests satisfy the monotone likelihood ratio property and the DM only wants to accept types above a threshold, the optimal menu is to only offer the most informative test.

Proposition 3. *Suppose $\Theta, X \subset \mathbb{R}$, $A = \{\theta : \theta > \overline{\theta}\}$ for some $\overline{\theta}$ and all tests in T have full-support and the monotone likelihood ratio property: for $\theta > \theta'$,*

$$\frac{\pi_t(x|\theta)}{\pi_t(x|\theta')} \text{ is increasing in } x.$$

If there is $t \succeq t'$ for all $t' \in T$, then, the menu $\{t\}$ is optimal.

Again this result holds by showing a single-crossing difference property on the acceptance probability. Intuitively, the reason is that more informative tests send relatively higher signals for higher types. So if a low type chooses the most informative test, the higher types must also choose that one. This prevents any pooling of A-types and R-types on two different tests. Combined with Lemma 1 that guarantees the inclusion of the dominant test, we get our result. Note also that this result would hold using weaker information order like Lehmann (1988) or some weakening of it. The key property delivering the result is the single-crossing condition described above.

If it is possible to find two signals, x, x', two A-types A_1, A_2 and one R-type, R_1 such that $\frac{\pi_t(x|A_1)}{\pi_t(x'|A_1)} < \frac{\pi_t(x|R_1)}{\pi_t(x'|R_1)} < \frac{\pi_t(x|A_2)}{\pi_t(x'|A_2)}$, then there is a test t' strictly less informative than t and a prior such that offering $\{t, t'\}$ is better for the DM than just offering $\{t\}$.

Proposition 4. *Let t be a test. Suppose there are two signals $x, x' \in X$, types $A_1, A_2 \in A$ and $R_1 \in R$ such that*

$$\frac{\pi_t(x|A_1)}{\pi_t(x'|A_1)} < \frac{\pi_t(x|R_1)}{\pi_t(x'|R_1)} < \frac{\pi_t(x|A_2)}{\pi_t(x'|A_2)}.$$

There is a prior μ and a test $t' \prec t$ such that the DM's payoffs are higher in the menu $\{t, t'\}$ than in $\{t\}$.

Intuitively, if we interpret x as a high signal, the A-type A_1 sends relatively low signals. Suppose that the prior is such that, if only t is offered, x is accepted and x' is not. In a sense, it means that in the test t, type R_1 performing better than A_1 on the signals x, x'. It is then beneficial for the DM to include a test that pools signals x, x' together. In that new test, type A_1 can choose the coarsened test where the superior performance of type R_1 is less important than in the original test.

The proof of Proposition 4 actually uses the following criterion to determine whether a less informative is part of the optimal menu. It gives condition to include coarsened version of a test.

Definition 3. *A test t is a coarsening of test t' if there is a partition of X, $\{X_i\}$, such that for all $\theta \in \Theta$,*

$$\pi_t(x_i|\theta) = \sum_{x \in X_i} \pi_{t'}(x|\theta) \quad \textit{for some } x_i \in X_i$$

$$\pi_t(x'|\theta) = 0 \quad \textit{for all } x' \in X_i,\, x' \neq x_i$$

The idea of a coarsening is that it pools all the signal in one element of the partition X_i on one signal x_i. The test t' is more informative than t as any strategy under t can be implemented under t'. I say that a test pools signals in X' if the partition is $\{X', \{x\} : x \notin X'\}$. Let $z^+ = \max\{0, z\}$.

Proposition 5. *Let $\alpha(x,t)$ be the optimal strategy when only test t is used. If there is $\tilde{\alpha} \in [0,1]$ and $X' \subseteq X$ such that*

$$\sum_{\theta \in A} \sum_{x \in X'} \mu(\theta) \big[(\tilde{\alpha} - \alpha(x,t))\pi_t(x|\theta)\big]^+ \geq \sum_{\theta' \in R} \sum_{x \in X'} \mu(\theta') \big[(\tilde{\alpha} - \alpha(x,t))\pi_t(x|\theta')\big]^+$$

then it is optimal to include a coarsened version of t that pools signals in X'.

This result is a direct application of the max-min problem of Theorem 1. It considers using the same strategy as in test t for the coarsened test but for the coarsened signal in X' where it uses $\tilde{\alpha}$. The condition then boils down to checking for a profitable deviation. The intuition for Proposition 5 is the same as in Proposition 4. The set X' identifies a set of signals where some A-types are performing worse than R-types. Offering a test that coarsens signals in X' creates a profitable deviation for these A-types.

4.2 Optimal menu with tests ordered by their difficulty

In many economic environments, the DM does not necessarily have access to a most informative test but can vary the difficulty to pass a test. This is for example the case for a regulator that can decide how demanding a certification test is. Like in Proposition 1 and Proposition 3, I show that the optimal menu is a singleton.

I first formalise the notion of more difficult test as follows.

Definition 4 (Difficulty environment)**.** *An environment is a Difficulty environment if* $\Theta \in \mathbb{R}$, $A = \{\theta : \theta > \overline{\theta}\}$ *for some* $\overline{\theta}$, $X = \{x_0, x_1\}$, $T \subset \mathbb{R}$, *all tests have full-support, satisfy the monotone likelihood ratio property and for all* $t > t'$, *and* $\theta > \theta'$,

$$\frac{\pi_t(x_1|\theta)}{\pi_t(x_1|\theta')} \geq \frac{\pi_{t'}(x_1|\theta)}{\pi_{t'}(x_1|\theta')} \quad \textit{and} \quad \frac{\pi_t(x_0|\theta)}{\pi_t(x_0|\theta')} \geq \frac{\pi_{t'}(x_0|\theta)}{\pi_{t'}(x_0|\theta')}$$

If $t > t'$, I will say that t is harder than t'. To understand the last condition better, let $\mu(\cdot|x, t)$ be a posterior belief after observing signal x in test t. The monotone likelihood ratio property implies $\mu(\cdot|t, x_1) \succeq_{FOSD} \mu(\cdot|t, x_0)$, a higher signal is "good news" about the type (Milgrom, 1981). The last property in the definition further implies $\mu(\cdot|t, x) \succeq_{FOSD} \mu(\cdot|t', x)$. That means that a pass grade shifts beliefs more towards higher type in a harder test and a fail grade shifts more beliefs towards lower types in an easy test. Or put differently, the harder a test the more informative it is about a high type when there is a pass-grade whereas an easier test is informative about the low types when the test is failed. As an example, if $\Theta \subset (0, 1)$ and $\pi_t(x_1|\theta) = \theta^t$ we are in a Difficulty environment.

Proposition 6. *In a Difficulty environment, a singleton menu is optimal.*

Like Proposition 1 and Proposition 3, Proposition 6 illustrates how incentive constraints shape the size of the optimal menu. In the case of the single-peaked preferences with dominant test, the equilibrium when the most informative test is offered is unique and only that test is chosen. Here, it is possible to construct an equilibrium where more than one test is chosen in equilibrium. However, the DM strategy needed to sustain that equilibrium is such that he is better off offering only one test.

The proof proceeds in two steps. First, I show that there are at most two tests in the optimal menu and if there are two tests, the harder test must be more lenient that the easy test. In particular, I show that after the hard test, the DM must accept with some probability after a fail signal and in the easy test, reject with positive probability after a pass grade.

This means that to maintain incentives to select both tests, the DM only reacts to the least informative signal from the test: in the hard test after a fail grade, in the easy test after a pass grade. This in turn implies that it would be better for the DM to use only one test and reject after a fail grade and accept after a pass grade.

4.3 Bidimensional environment

In this subsection, I apply the tools of Theorem 1 to study environments with bidimensional types. The analysis here can be easily extended to more than two dimensions. I assume that the DM has access to tests that can only reveal one dimension and the preference of the DM have some monotonicity along each dimension.

Definition 5. *An environment is bidimensional if* $\Theta = \Theta_1 \times \Theta_2 \subset \mathbb{R}^2$*,* $X \subset \mathbb{R}$ *and* $T = \{t_1, t_2\}$ *such that for* $i = 1, 2$*,*

- *if* $\theta \in A$*, then for all* $\theta' \geq \theta$*,* $\theta' \in A$

- *t_i has full support and for all $\theta_i > \theta_i'$,*

$$\frac{\pi_{t_i}(x|\theta_i,\theta_j)}{\pi_{t_i}(x|\theta_i',\theta_j)} \text{ is strictly increasing in } x \text{ for any } \theta_j \in \Theta_j$$

- *$\pi_{t_i}(x|\theta_i,\theta_j) = \pi_{t_i}(x|\theta_i,\theta_j')$ for all $\theta_j, \theta_j' \in \Theta_j$ and $x \in X$*

The first condition captures the idea that a higher type is always better for the DM. The second and third condition captures the idea that each test is only informative about one dimension and that a higher signal corresponds to a higher type in that dimension.

In this environment, whether the DM wants to offer a menu depends crucially on his preferences. In particular, I give a necessary and sufficient condition on the preferences such that a menu is optimal for any prior. Let $\overline{\theta}_i = \max \Theta_i$.

Proposition 7. *Suppose we are in a bidimensional environment. Offering a menu $\{t_1, t_2\}$ is strictly optimal for any prior if and only if*

$$\text{for } i = 1, 2,\ (\overline{\theta}_i, \theta_j) \in A, \text{ for all } \theta_j \in \Theta_j. \tag{2}$$

The proof of Proposition 7 works by showing that a deviation from a single test menu is always profitable when condition (2) is satisfied and constructs a prior under which there are no profitable deviations when the condition is not satisfied.

Figure 1 illustrates the condition of Proposition 7 with $\Theta \subset [0,1]^2$. In Figure 1a, the DM wants the agent's type to be high enough in at least one dimension. Then the DM always prefers to offer a full menu to the agent. On the other hand, in Figure 1b, the DM does not want to accept a type that is high in only one dimension. In this case, for some prior, the DM only wants to offer one test. This happens when after any deviation from the singleton menu any A-type is mimicked by too many R-types that cannot be distinguished from him.

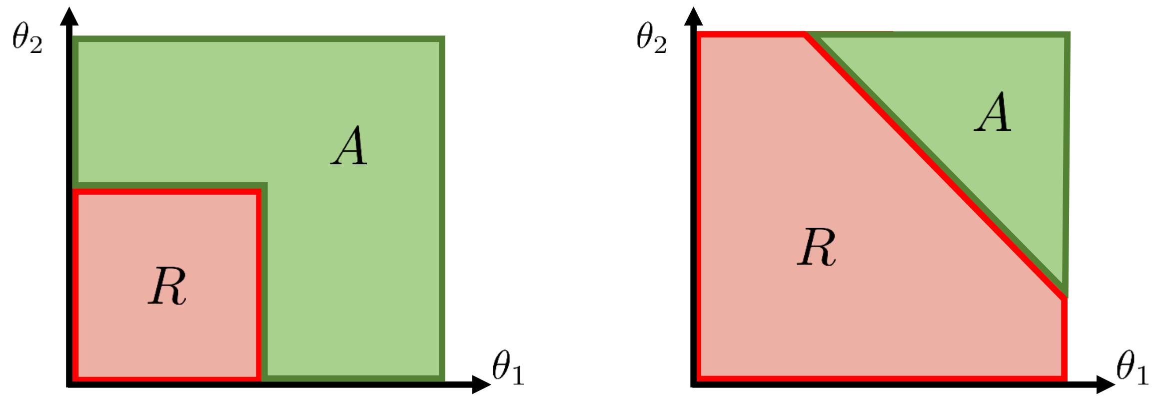

(a) Full menu always optimal. (b) Full menu not always optimal.

Figure 1: Illustration of DM's preferences for Proposition 7.

5 Sufficient conditions for test inclusion

In this section, I study in more details the notion of efficient allocation of tests to the agent's types. I show that a sufficient condition to include a test in the optimal menu is if it is good at differentiating one A-type from all the R-types. This captures a notion of a test tailored for the A-type.

Definition 6. *Fix $\theta \in A$. Test t θ-dominates t', $t \succeq_\theta t'$, if there is $\beta : X \times X \to [0,1]$ such that for all $x' \in X$*

$$\begin{aligned} & \sum_x \beta(x,x')\pi_t(x|\theta) \leq \pi_{t'}(x'|\theta) \\ \textit{for all } \theta' \in R, \quad & \sum_x \beta(x,x')\pi_t(x|\theta') \geq \pi_{t'}(x'|\theta') \\ \textit{for all } x \in X, \quad & \sum_{x'} \beta(x,x') \leq 1 \end{aligned}$$

To understand this definition better, compare it to Blackwell (1953)'s informativeness order. It requires the existence of a function β such that for all $x' \in X$, $\sum_x \beta(x,x')\pi_t(x|\theta) = \pi_{t'}(x'|\theta)$ for all $\theta \in \Theta$ and for all $x \in X$, $\sum_{x'} \beta(x,x') = 1$. The key difference is that we restrict attention to one A-type and all the R-types. This captures the idea the test θ-dominant test is tailored to differentiate θ from each R-type. The second difference is that it requires only inequalities whereas the Blackwell order requires equalities. This is because we are fixing the utility function we are interested in, unlike in Blackwell (1953).

If a type $\theta \in A$ has a $\succeq_\theta$-dominant test, then this test is used in an optimal menu. This shows that an important property of tests is not so much how good they are at differentiating types, but how good they are at differentiating one type the DM wants to accept from all the types he wants to reject.

Proposition 8. *Suppose there is $t \in T$ and $\theta \in A$ such that $t \succeq_\theta t'$ for all $t' \in T$, then t is part of an optimal menu.*

The stronger notion of a test able to differentiate some $\theta \in A$ from all R-types is if $\operatorname{supp} \pi_t(\cdot|\theta) \cap \left(\cup_{\theta' \in R} \operatorname{supp} \pi_t(\cdot|\theta') \right) = \emptyset$. If each type in A has such a test, then the principal never makes a mistake. This condition is also necessary.

Proposition 9. *The principal's expected payoff is $\sum_{\theta \in A} \mu(\theta)$ if and only if for all $\theta \in A$, there exists $t \in T$ such that*

$$\operatorname{supp} \pi_t(\cdot|\theta) \cap \left(\cup_{\theta' \in R} \operatorname{supp} \pi_t(\cdot|\theta') \right) = \emptyset$$

Here, the principal just needs for each type he wants to accept a test where he can discriminate between that type and the R-types. Then he can offer a menu of tests where each A-type self selects into the test that discriminates him from the R-types. The actual learning only happens by observing the test selected and the testing technology serves as a detriment to deviations from R-types. The argument is then similar to an unravelling argument à la Milgrom (1981) and Grossman (1981). These are not fully revealing tests but tests that allow to perfectly discriminate *one* A-type from all the R-types. But it could be a very noisy tests for the other A-types.

6 Extension: Communication

I consider here the possibility of adding a communication channel on top of the test choice. There is now a finite set C of output messages with $|C| \geq |A|$ and a strategy is a mapping $\sigma : \Theta \to \Delta(T \times C)$. Note that all the results from the previous sections go through as from any finite set T one can create another T' that duplicate each test $|C|$ times. I call this variant of the model the *menu game with communication*.

In line with Theorem 1, each A-type chooses a message-test pair deterministically and each R-type mixes over some A-types message-test pair. Moreover, I show that when communication is added, each type in A announces his type, thus maximally differentiating himself, and each R-type pretends to be an A-type.

Theorem 2. *If communication is allowed, the same construction as Theorem 1 holds. Moreover, there is a DM-preferred equilibrium where each A-type reports his own type.*

Proof. See appendix. □

Theorem 2 shows that the results extend naturally to an environment where communication is allowed. Because the DM could commit to a strategy, he can always guarantee each A-type at least as much as he would have if he would pool with another A-type. This guarantees that there is an a solution to the max-min problem where he separates from the other A-types.

Note that because each A-type uses a different message and does not mix over tests, the test chosen does not contain any information: $\mu(\theta|c,t) = \mu(\theta|c)$. Thus all the information revealed by the test is through the signal realisations and not the test choice.

In the remainder of this section, I will connect the results developed in this model to the existing literature, and in particular to Glazer and Rubinstein (2004). Consider the following model generalising the one of Glazer and Rubinstein (2004). They consider a model of persuasion and verification where the agent sends a message and the DM chooses a test and a decision based on the message. Formally, the DM designs a mechanism defined by $\tau : C \to \Delta(T \times [0,1]^X)$, that is a mechanism commits to a test and a decision for each test and signal realisation for each message. A strategy for the agent is $\delta : \Theta \to \Delta C$. The solution concept is Perfect Bayesian Equilibrium. In Glazer and Rubinstein (2004), the state space is some multidimensional set and each test in T perfectly reveals one dimension. I will call the mechanism τ a GR-mechanism.

One of the results of Glazer and Rubinstein (2004) is that the outcome of the optimal mechanism τ can be implemented without commitment in a PBE of the following game: the agent chooses a message in C, based on the message, the DM chooses a test and based on the signal realisation and test, the DM accepts or rejects the agent. If the outcome of the optimal GR-mechanism is the same as the one of the game above, I will say that it is credible. I will call that game a GR-game.

The fundamental difference between the Glazer and Rubinstein (2004) model and the one we have studied so far is that it is now the DM that chooses the test and not the agent. But as we will see, if we allow for communication in the menu of test model, this distinction does not matter anymore.

Proposition 10. *The outcome of an optimal GR-mechanism is credible for any T. Moreover, its outcome coincides with the optimal menu game with communication.*

This proposition generalises the commitment result of Glazer and Rubinstein (2004) to an arbitrary testing technology and type structure.[5] Moreover, it shows that when there is communication, who chooses the test is not important. To understand this better, let us first note the dual role of test choice in the model without communication. In this case, the test is used both to communicate to the DM and to provide evidence which type the agent is. When we add communication on top of the menu of test, all the communication is through the cheap-talk message and the test is only used to provide evidence about the type.

Now consider the max-min problem characterisation of the optimal menu, and in particular the payoffs of the A-types. Remember that in the max-min problem, the A-types were maximising the DM's payoffs. Combined with the fact that the test choice does not carry additional information, we can let the DM choose it. If it was optimal for A-types to choose test t after message c, it will also be for the DM.

Carroll and Egorov (2019) study a similar model as Glazer and Rubinstein (2004), multidimensional types with the testing technology revealing one dimension, but with a different agent payoff function. They study under which condition on the agent's payoffs there is full information revelation. They show that when there is full information revelation and some technical conditions are satisfied, the mechanism can be implemented by having the agent choosing the test, a parallel result to Proposition 10. Thus I show that the equivalence result they have also applies to other environments and is not a feature of full information revelation and their testing technology.

[5] Weksler and Zik (2022) and Silva (2022) study a related model where the agent sends a report to the DM who has private information about the agent. This can be thought as having a unique test and communication from the agent.

7 Conclusion

I study the design of optimal menus of tests. Menus allow the DM to have an additional dimension for information revelation as well as allow for a more efficient allocation of tests to the agent's types. I provide a characterisation of the optimal menu in terms of an auxiliary max-min problem. One advantage of this characterisation is that it does not rely on any structure on types or tests. While proving this result, I also show that the characterisation holds for a general class of mechanisms allocating agent to tests.

In applications, I show that using a menu can be a powerful tool, and even a dominated test, in the Blackwell sense, can be part of the optimal menu. However, this channel also has limits and I show that in some natural economic environments the optimal menu is a singleton. All the results also hold when the DM can commit to an action. I interpreted this result as a hierarchy over information sources: even when the DM can use a suboptimal strategy to "artificially" incentivise the agent to choose different tests, he is better off using a menu only when he can best reply to the information revealed.

Results for the optimality of the inclusion of some tests, like Proposition 5 and Proposition 8, reveal an interesting asymmetry between types. They are comparing properties of a test or acceptance probability of one or some A-type to those of all the R-types. This asymmetry between the R-types and the A-types is due to their different incentives to separate as reflected by their strategy in the auxiliary max-min problem. While an A-type wants to be singled-out by choosing a different test, the R-types want to "hide behind" A-types and only choose to mimic them. Proposition 5 and Proposition 8 provide conditions under which an A-type is better off deviating to a new test while providing limited incentives to the R-types to mimic him.

Finally, I show that adding a communication channel links the current model to existing models in the literature and generalises their results. Adding the communication highlights the role of tests when there is no communication. Without communication, the tests also serve as a communication channel. When communication is allowed, the test choice does not add any information beyond the test results. The DM is thus as well off choosing the test himself following the cheap-talk message.

References

Ball, I. (2021), 'Scoring strategic agents', *arXiv preprint arXiv:1909.01888* .

Ball, I. and Kattwinkel, D. (2022), 'Probabilistic verification in mechanism design'.

Ben-Porath, E., Dekel, E. and Lipman, B. L. (2019), 'Mechanisms with evidence: Commitment and robustness', *Econometrica* **87**(2), 529–566.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA14991*

Ben-Porath, E., Dekel, E., Lipman, B. L. and Draft, F. P. (2021), 'Mechanism design for acquisition of/stochastic evidence'.

Blackwell, D. (1953), 'Equivalent comparisons of experiments', *The Annals of Mathematical Statistics* **24**(2), 265–272.

Boleslavsky, R. and Kim, K. (2018), 'Bayesian persuasion and moral hazard', *Available at SSRN 2913669* .

Bull, J. and Watson, J. (2007), 'Hard evidence and mechanism design', *Games and Economic Behavior* **58**(1), 75–93.

Carroll, G. and Egorov, G. (2019), 'Strategic communication with minimal verification', *Econometrica* **87**(6), 1867–1892.

Deb, R. and Stewart, C. (2018), 'Optimal adaptive testing: informativeness and incentives', *Theoretical Economics* **13**(3), 1233–1274.
URL: *https://econtheory.org/ojs/index.php/te/article/view/2914/0*

DeMarzo, P. M., Kremer, I. and Skrzypacz, A. (2019), 'Test design and minimum standards', *American Economic Review* **109**(6), 2173–2207.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20171722*

Deneckere, R. and Severinov, S. (2008), 'Mechanism design with partial state verifiability', *Games and Economic Behavior* **64**(2), 487–513.

Dye, R. A. (1985), 'Disclosure of nonproprietary information', *Journal of accounting research* pp. 123–145.

Ely, J., Galeotti, A., Jann, O. and Steiner, J. (2021), 'Optimal test allocation', *Journal of Economic Theory* **193**, 105236.

Forges, F. and Koessler, F. (2005), 'Communication equilibria with partially verifiable types', *Journal of Mathematical Economics* **41**(7), 793–811.

Giovannoni, F. and Seidmann, D. J. (2007), 'Secrecy, two-sided bias and the value of evidence', *Games and Economic Behavior* **59**(2), 296–315.
URL: *https://www.sciencedirect.com/science/article/pii/S0899825606001047*

Glazer, J. and Rubinstein, A. (2004), 'On optimal rules of persuasion', *Econometrica* **72**(6), 1715–1736.

Glazer, J. and Rubinstein, A. (2006), 'A study in the pragmatics of persuasion: a game theoretical approach', *Theoretical Economics* **1**(4), 395–410.
URL: *https://econtheory.org/ojs/index.php/te/article/view/20060395/0*

Green, J. R. and Laffont, J.-J. (1986), 'Partially verifiable information and mechanism design', *The Review of Economic Studies* **53**(3), 447–456.

Grossman, S. J. (1981), 'The informational role of warranties and private disclosure about product quality', *The Journal of Law and Economics* **24**(3), 461–483.

Hagenbach, J., Koessler, F. and Perez-Richet, E. (2014), 'Certifiable pre-play communication: Full disclosure', *Econometrica* **82**(3), 1093–1131.

Harbaugh, R. and Rasmusen, E. (2018), 'Coarse grades: Informing the public by withholding information', *American Economic Journal: Microeconomics* **10**(1), 210–35.
URL: *https://www.aeaweb.org/articles?id=10.1257/mic.20130078*

Hart, S., Kremer, I. and Perry, M. (2017), 'Evidence games: Truth and commitment', *American Economic Review* **107**(3), 690–713.

Hubler, S. (2020), 'Why Is the SAT Falling Out of Favor?', *The New York Times* .
URL: *https://www.nytimes.com/2020/05/23/us/SAT-ACT-abolish-debate-california.html*

Kartik, N. and Tercieux, O. (2012), 'Implementation with evidence', *Theoretical Economics* **7**(2), 323–355.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/TE723*

Koessler, F. and Perez-Richet, E. (2019), 'Evidence reading mechanisms', *Social Choice and Welfare* **53**(3), 375–397.

Koessler, F. and Skreta, V. (2022), 'Informed information design'.

Kreps, D. M. and Wilson, R. (1982), 'Sequential equilibria', *Econometrica* pp. 863–894.

Lehmann, E. L. (1988), 'Comparing location experiments', *The Annals of Statistics* **16**(2), 521–533.

Lipman, B. L. and Seppi, D. J. (1995), 'Robust inference in communication games with partial provability', *Journal of Economic Theory* **66**(2), 370–405.

Lipnowski, E., Ravid, D. and Shishkin, D. (2022), 'Persuasion via weak institutions', *Journal of Political Economy* **130**(10), 2705–2730.

Milgrom, P. (2008), 'What the seller won't tell you: Persuasion and disclosure in markets', *Journal of Economic Perspectives* **22**(2), 115–131.

Milgrom, P. R. (1981), ‘Good news and bad news: Representation theorems and applications’, *The Bell Journal of Economics* pp. 380–391.

Nguyen, A. and Tan, T. Y. (2021), ‘Bayesian persuasion with costly messages’, *Journal of Economic Theory* **193**, 105212.

Perez-Richet, E. and Skreta, V. (2022), ‘Test design under falsification’, *Econometrica* **90**(3), 1109–1142.

Perez-Richet, E., Vigier, A. and Bizzotto, J. (2020), ‘Information design with agency’.

Persico, N. (2000), ‘Information acquisition in auctions’, *Econometrica* **68**(1), 135–148.

Rockafellar, R. T. (2015), *Convex Analysis*, Princeton University Press.
URL: *https://doi.org/10.1515/9781400873173*

Rosar, F. (2017), ‘Test design under voluntary participation’, *Games and Economic Behavior* **104**, 632–655.

Sher, I. (2011), ‘Credibility and determinism in a game of persuasion’, *Games and Economic Behavior* **71**(2), 409–419.
URL: *https://www.sciencedirect.com/science/article/pii/S0899825610000953*

Silva, F. (2022), ‘Information transmission in persuasion models with imperfect verification’, *Available at SSRN 3728325* .

Strausz, R. (2017), Mechanism Design with Partially Verifiable Information, Rationality and Competition Discussion Paper Series 45, CRC TRR 190 Rationality and Competition.
URL: *https://ideas.repec.org/p/rco/dpaper/45.html*

Weksler, R. and Zik, B. (2022), ‘Job market cheap talk’.

A Omitted proofs

A.1 Proof of Theorem 1

The plan of the proof is the following. First, I characterise the optimal mechanism, where a mechanism maps an input message to a distribution over tests. Because an equilibrium of the menu game can be implemented by a mechanism, the payoffs from the optimal mechanism are weakly greater than the payoffs from any optimal menu. In the second part of the proof, I show that the optimal mechanism can be implemented by posting a menu. In the proof, I will refer to a distribution over test as an allocation.

By standard arguments, a direct truthful mechanism is without loss of generality. A direct mechanism is a mapping $\tilde{\sigma} : \Theta \to \Delta T$. The designer's problem is

$$\begin{aligned} \tilde{V} = \max_{\tilde{\sigma},\alpha} & \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) - \sum_{\theta \in R} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) \\ \text{s.t.} & \sum_t (\tilde{\sigma}(t|\theta) - \tilde{\sigma}(t|\theta')) p_t(\alpha;\theta) \geq 0 \text{ for all } \theta, \theta' \\ & \alpha \in BR(\tilde{\sigma}) \end{aligned}$$

The first constraint is the incentive compatibility constraint of type θ deviating to θ', the second guarantees that an allocation is well-defined and the last constraint ensures that the DM best replies to the information revealed by the output of the mechanism.

Note that any equilibrium in the menu game is incentive compatible and therefore a solution to $\tilde{V}$ gives weakly higher expected payoffs to the DM.

If the DM could commit over a strategy α, his problem would be

$$\tilde{V}(\alpha) = \max_{\tilde{\sigma}} \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) - \sum_{\theta \in R} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta)$$
$$\text{s.t. } \sum_t (\tilde{\sigma}(t|\theta) - \tilde{\sigma}(t|\theta')) p_t(\alpha;\theta) \geq 0 \text{ for all } \theta, \theta'$$

We have that $\max_\alpha V(\tilde{\alpha}) \geq \tilde{V}$ as the DM could always commit to the strategy used to get $\tilde{V}$.

Step 1: Show that $\tilde{V}(\alpha) = \max_s \min_m v(\alpha, s, m)$ where v is defined in (1).

To show this claim, I am going to relax the mechanism design problem by restricting attention to the IC constraints of R-types deviating to reporting an A-type:

$$\tilde{V}(\alpha) = \max_{\tilde{\sigma}} \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) - \sum_{\theta \in R} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta)$$
$$\text{s.t. } \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) \geq \max_{m(\cdot|\theta)} \sum_{\theta' \in A} m(\theta'|\theta) \sum_t \tilde{\sigma}(t|\theta') p_t(\alpha;\theta), \text{ for all } \theta \in R$$

The IC constraints are written to express that reporting type θ for $\theta \in R$ is better than any other reporting strategy over the A-types.

Now note that if an IC constraint is slack at the optimum, we could improve the DM's payoff by setting $\tilde{\sigma}(\cdot|\theta) = \max_{m(\cdot|\theta)} \sum_{\theta' \in A} m(\theta'|\theta) \sum_t \tilde{\sigma}(t|\theta') p_t(\alpha;\theta)$. As that would reduce the probability of type $\theta \in R$ of being accepted and would not change any other constraints in the relaxed problem. Thus at the optimum, $\tilde{\sigma}(\cdot|\theta) = \max_{m(\cdot|\theta)} \sum_{\theta' \in A} m(\theta'|\theta) \sum_t \tilde{\sigma}(t|\theta') p_t(\alpha;\theta)$. We can plug this expression in the payoffs to get

$$\tilde{V}(\alpha) = \max_{\tilde{\sigma}} \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) - \sum_{\theta \in R} \mu(\theta) \max_{m(\cdot|\theta)} \sum_{\theta' \in A} m(\theta'|\theta) \sum_t \tilde{\sigma}(t|\theta') p_t(\alpha;\theta)$$
$$= \max_{\tilde{\sigma}} \min_m \sum_{\theta \in A} \mu(\theta) \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha;\theta) - \sum_{\theta \in R} \mu(\theta) \sum_{\theta' \in A} m(\theta'|\theta) \sum_t \tilde{\sigma}(t|\theta') p_t(\alpha;\theta)$$

Note that we can take out the max of the summation by the linearity of the expression in m and it becomes a min because of the minus sign. This expression also corresponds to v as defined in (1).

It remains to show that the solution of this relaxed mechanism is indeed optimal. Take $s \in \arg\max \min_{\tilde{m}} v(\alpha, \tilde{s}, \tilde{m})$ and $m \in \arg\min v(\alpha, s, \tilde{m})$ and define the optimal mechanism by

- $\tilde{\sigma}(t|\theta) = s(t|\theta)$ for $\theta \in A$
- $\tilde{\sigma}(t|\theta') = \sum_{\theta \in A} m(\theta|\theta')s(t|\theta)$ for $\theta' \in R$

Note that an outcome of this mechanism gives payoff weakly higher than $\max_{\tilde{s}} \min_{\tilde{m}} v(\alpha, \tilde{s}, \tilde{m})$ as $v(\alpha, s, m) \geq \min_{\tilde{m}} v(\alpha, s, \tilde{m})$. Thus if it is incentive-compatible, it must be actually equal to the upper bound $\max_{\tilde{s}} \min_{\tilde{m}} v(\alpha, \tilde{s}, \tilde{m})$.

Note that all allocations are either allocations of A-types or convex combinations of the A-types' allocations thus it is enough to check that no type has any incentive to report any A-type in the mechanism.

By definition of m,

$$v(\alpha, s, m) \geq v(\alpha, s, m'), \text{ for all } m'$$

Therefore, for each R-type, θ',

$$\sum_{\theta \in A} m(\theta|\theta') \sum_t \tilde{\sigma}(t|\theta) p_t(\alpha; \theta') \geq \sum_t \tilde{\sigma}(t|\tilde{\theta}) p_t(\alpha; \theta'), \text{ for any } \tilde{\theta} \in A$$

For an A-type θ, consider the choice of choosing $\tilde{s}$ such as $\tilde{s}(\cdot|\theta) = s(\cdot|\tilde{\theta})$ for some $\tilde{\theta} \in A$ and the same otherwise. By definition of s,

$$v(\alpha, s, m) = \min_{\tilde{m}} v(\alpha, s, \tilde{m}) \geq \min_{\tilde{m}} v(\alpha, \tilde{s}, \tilde{m})$$

Rearranging,

$$\mu(\theta)\sum_t \big(s(t|\theta)-s(t|\tilde{\theta})\big)p_t(\alpha;\theta) \geq \max_{\tilde{m}}\sum_{\theta'\in R}\mu(\theta')\sum_{\theta''\in A}\tilde{m}(\theta''|\theta')\sum_t s(t|\theta)p_t(\alpha;\theta')$$
$$-\max_{\tilde{m}}\sum_{\theta'\in R}\mu(\theta')\sum_{\theta''\in A}\tilde{m}(\theta''|\theta')\sum_t \tilde{s}(t|\theta)p_t(\alpha;\theta')$$

where the min is transformed in max because of the negative sign. Note that the LHS is the IC constrain of type θ deviating to type $\tilde{\theta}$. The RHS is the difference payoff is the probability of the R-types of being accepted when they choose a mimicking strategy $\tilde{m}$. Note that the only difference between s and $\tilde{s}$, from their point of view is that there is weakly less choice of allocations to mimic as we have θ choosing the same allocation as $\tilde{\theta}$. Therefore it must be that the RHS is positive which implies that the LHS is as well.

Step 2: Show that the DM does not benefit from commitment in the optimal mechanism.

Take $(\alpha, s) \in \arg\max_{\tilde{\alpha},\tilde{s}} \min_{\tilde{m}} v(\tilde{\alpha},\tilde{s},\tilde{m})$ and $m \in \arg\min_{\tilde{m}}\max_{\tilde{\alpha}} v(\tilde{\alpha}, s, \tilde{m})$. The α selected would be the optimal strategy when the DM can commit.

Note that because the order of maximisation does not matter, we also have $\alpha \in \arg\max_{\tilde{\alpha}}\min_{\tilde{m}} v(\tilde{\alpha}, s, \tilde{m})$. Note that v is linear in $\tilde{\alpha}$ and $\tilde{m}$ and thus by the minimax theorem,

$$v(\alpha, s, m) \geq v(\alpha', s, m), \text{ for all } \alpha'$$
$$v(\alpha, s, m) \leq v(\alpha, s, m'), \text{ for all } m'$$

Thus α best-replies to the optimal mechanism when the DM can commit and m is also a best reply to (α, s), thus satisfying the condition for characterising the equilibrium in Step 1.

Step 3: Show that an optimal mechanism can be implemented by posting a menu.

Lemma 2. *Take* $(\alpha, s) \in \arg\max_{\tilde{\alpha},\tilde{s}} \min_{\tilde{m}} v(\tilde{\alpha},\tilde{s},\tilde{m})$ *and* $m \in \arg\min_{\tilde{m}}\max_{\tilde{\alpha}} v(\tilde{\alpha}, s, \tilde{m})$.

If $s(\cdot|\theta)$ is in pure strategy for all $\theta \in A$, then the optimal mechanism $\tilde{\sigma}$ with DM strategy α is implementable by posting a menu where the strategies are

- *for $\theta \in A : \sigma(t|\theta) = s(t|\theta)$*
- *for $\theta' \in R : \sigma(t|\theta') = \sum_{\theta \in A} m(\theta|\theta')s(t|\theta)$*
- *the DM's strategy is α.*

Moreover, the DM does not benefit from committing to α.

Proof. Note that the strategies σ are the same as the outcome of the optimal mechanism $\tilde{\sigma}$ when the DM strategy is α.

Optimal mechanism is implementable with a menu.

The menu posted by the DM is $\mathcal{M} = \cup_{\theta \in A} \operatorname{supp} s(\cdot|\theta)$. To prove the result, we simply need to show that the pair (σ, α) is a PBE in the game when the menu $\mathcal{M}$ is posted. Let t^θ be the test chosen by type $\theta \in A$.

The incentive compatibility constraint of type $\theta \in A$ deviating to $\tilde{\theta} \in A$ in the optimal mechanism implies

$$p_{t^\theta}(\alpha; \theta) \geq p_{t^{\tilde{\theta}}}(\alpha; \theta)$$

for any $\tilde{\theta} \in A$. Thus $\theta \in A$ prefers t^θ to any other $t' \in \mathcal{M}$.

The incentive compatibility constraint of type $\theta' \in R$ deviating to $\tilde{\theta} \in A$ in the optimal mechanism implies

$$\sum_t \sigma(t|\theta')p_t(\alpha; \theta') = \sum_t \sum_{\theta \in A} m(\theta|\theta')s(t|\theta)p_t(\alpha; \theta') \geq p_{t^{\tilde{\theta}}}(\alpha; \theta)$$

which again implies that $\sum_t \sigma(t|\theta')p_t(\alpha; \theta') \geq p_{t'}(\alpha; \theta')$ for all $t' \in \mathcal{M}$.

In Step 2, we have shown that the strategy is a best-reply to σ.

On-path, beliefs are pinned down by the strategy σ, the tests π_t and the prior. Off-path, we can choose a belief $\tilde{\mu}(\cdot|t,x)$ such that $\alpha(t,x)$ is a best-reply to $\tilde{\mu}$.

This conclude the description of the PBE.

No benefit to commitment.

This follows from the fact that the DM does not benefit from commitment in the optimal mechanism and that the payoffs from the optimal menu without commitment are the same as in the optimal mechanism with commitment. Given that the payoffs from the optimal mechanism with commitment are always weakly higher than the optimal menu with commitment, the DM does not benefit from commitment to α in the optimal menu. $\square$

Lemma 3. *For any α, there is $s^* \in \arg\max_s \min_m v(\alpha, s, m)$ such that $s^*(\cdot|\theta)$ is in pure strategy for all $\theta \in A$.*

Proof. Suppose there is $s^* \in \arg\max_s \min_m v(\alpha, s, m)$ such that for some $\theta \in A$, and $t, t' \in T$, $s^*(t|\theta), s^*(t'|\theta) > 0$.

Assume first that for any t, t', and $Z \subseteq R$,

$$\mu(\theta)p_t(\alpha;\theta) - \sum_{\theta' \in Z} \mu(\theta')p_t(\alpha;\theta') \neq \mu(\theta)p_{t'}(\alpha;\theta) - \sum_{\theta' \in Z} \mu(\theta')p_{t'}(\alpha;\theta') \tag{3}$$

Note that if $s^* \in \arg\max_s \min_m v(\alpha, s, m)$, it has to be optimal for any selection of $\arg\min_m v(\alpha, s, m)$. Take the selection, for all $\theta' \in R$,

$$m(\theta|\theta') = 1 \Leftrightarrow \sum_t s(t|\theta)p_t(\alpha;\theta') \geq \sum_t s(t|\tilde{\theta})p_t(\alpha;\theta'), \text{ for all } \tilde{\theta} \in A$$

that is, whenever mimicking θ is a best-reply for $\theta' \in R$, that type mimics θ with probability one. Because payoffs are linear, this is a best-reply. Now note that when θ evaluates his payoffs with respect to that selection, any $\theta' \in R$ that does not mimic him, strictly prefers another type. Moreover, by condition (3), type θ strictly benefits from putting $\epsilon > 0$ more weight on either t or t' for ϵ small enough. Indeed by changing the weight a little bit, he can increase his payoff and if ϵ is small enough no new type $\theta' \in R$ wants to mimic him. So this is a profitable deviation.

Now note that any payoffs satisfying condition (3) defines a dense subset of the payoff space, $(p_t(\alpha;\theta)_{t\in T,\theta\in\Theta})$, using the usual metric for $\mathbb{R}^n$. Indeed, condition (3) is a finite system of inequalities and any perturbation to $p_t(\alpha;\theta)$ upsets any equality. Take a sequence in the payoff space such that for any member of the sequence, condition (3) is satisfied such that the sequence converges to an element of the payoff space where condition (3) is not satisfied. Take an associated sequence of $s^{*,n} \in \arg\max_s \min_m v^n(\alpha, s, m)$ where n indexes the sequence. $(s^{*,n})$ is a bounded sequence in a closed subset of $\mathbb{R}^n$ so it admits a converging subsequence. This subsequence contains only pure strategies so it must converge to a pure strategy. By upper hemicontinuity of the Nash Equilibrium correspondence, the limit is a Nash Equilibrium and thus there is $s^* \in \arg\max_s \min_m v(\alpha, s, m)$ in pure strategy for any payoff. $\square$

A.2 Proof of Lemma 1

Because $t \succeq t'$ implies $t \succeq_\theta t'$ for some $\theta \in A$, Lemma 1 is a corollary of Proposition 8 proven below.

A.3 Proof of Proposition 1 and Proposition 2

Suppose the DM's preferences are single-peaked given $\geq_t$. Suppose there is a menu with both t, t'. Take $A_1, A_2 \in A$ with $A_1 < A_2$ and without loss of generality, suppose A_1 chooses t' and A_2 chooses t in some equilibrium. Let α denote the DM equilibrium strategy in this equilibrium.

Because $t \succeq t'$, there is $\beta : X \times X \to [0,1]$ such that $p_{t'}(\tilde{x}|\theta) = \beta(x,\tilde{x})\pi_t(x|\theta) + \beta(x',\tilde{x})\pi_t(x'|\theta)$ and $\sum_x \beta(\tilde{x},x) = 1$ for $\tilde{x} = x, x'$. Type $\theta \in \Theta$ prefers test t' over t if

$$\alpha(x_1,t')\Big(\beta(x_1,x_1)\pi_t(x_1|\theta) + \beta(x_0,x_1)\pi_t(x_0|\theta)\Big)$$
$$+\alpha(x_0,t')\Big(\beta(x_1,x_0)\pi_t(x_1|\theta)+\beta(x_0,x_0)\pi_t(x_0|\theta)\Big)-\alpha(x_1,t)\pi_t(x_1|\theta)-\alpha(x_0,t)\pi_t(x_0|\theta) \geq 0$$

Note that this expression is monotonic in θ. Indeed, if $\pi_t(x_0|\theta) > 0$, then dividing by $\pi_t(x_0|\theta)$ gives

$$\alpha(x_1,t')\Big(\beta(x_1,x_1)\frac{\pi_t(x_1|\theta)}{\pi_t(x_0|\theta)} + \beta(x_0,x_1)\Big) + \alpha(x_0,t')\Big(\beta(x_1,x_0)\frac{\pi_t(x_1|\theta)}{\pi_t(x_0|\theta)} + \beta(x_0,x_0)\Big)$$
$$- \alpha(x_1,t)\frac{\pi_t(x_1|\theta)}{\pi_t(x_0|\theta)} - \alpha(x_0,t)$$

which is linear in $\frac{\pi_t(x_1|\theta)}{\pi_t(x_0|\theta)}$, an increasing function of θ. If $\pi_t(x_0|\theta) = 0$, then $\pi_t(x_0|\theta') = 0$ for all $\theta' >_t \theta$ and the expression is constant.

To have A_1 choose t' and A_2 choose t, it must be strictly decreasing[6] in θ, i.e.,

$$\alpha(x_1,t')\beta(x_1,x_1) + \alpha(x_0,t')\beta(x_1,x_0) - \alpha(x_1,t) < 0 \tag{4}$$

[6] If all types are indifferent between t and t' then it is also an equilibrium to offer only t and the DM's payoffs are the same.

A necessary condition for (4) to hold is that $\alpha(x_1,t) > 0$. Note the strict monotonicity also implies that there is $\overline{\theta} \in A$ such that any $\theta > \overline{\theta}$ prefers t and any $\theta \leq \overline{\theta}$ prefers t'. Let $A^+ = \{\theta \in A : \theta >_t \overline{\theta}\}$ and $R^+ = \{\theta \in R : \theta >_t \theta',\ \text{for all } \theta' \in A\}$. But because only types in $A^+ \cup R^+$ choose t, the likelihood ratios $\frac{\pi_t(x_1|\theta)}{\pi_t(x_1|\theta')} < \frac{\pi_t(x_0|\theta)}{\pi_t(x_0|\theta')}$ for any $\theta \in A^+$, $\theta' \in R^+$ and $\alpha(x_1,t) > 0$ imply that $\alpha(x_0,t) = 1$ (Milgrom, 1981).

But then no type ever prefer t' over t. Indeed, the condition to prefer t' over t,

$$\begin{aligned}\Big(\alpha(x_1,t')\beta(x_1,x_1) + \alpha(x_0,t')\beta(x_1,x_0) - \alpha(x_1,t)\Big)&\pi_t(x_1|\theta)\\ \geq \Big(1 - \alpha(x_1,t')\beta(x_0,x_1) - \alpha(x_0,t')\beta(x_0,x_0)\Big)&\pi_t(x_0|\theta)\end{aligned}$$

is never satisfied as the LHS is strictly negative because (4) must hold and the RHS is positive because $\beta(x_0,x_1) + \beta(x_0,x_0) = 1$ and $\alpha(\tilde{x},t') \leq 1$, $\tilde{x} = x_1, x_0$.

Thus there cannot be an equilibrium where another test than t is chosen.

Suppose the DM's preferences are enclosed given $\geq_t$.

Suppose $(\tilde{\alpha},\tilde{s}) \in \arg\max \min_m v(\alpha,s,m)$ with $\tilde{s}(t|\theta) = 1$ for all $\theta \in A$.

Suppose the prior is such that when only t is offered, x_0 is rejected and x_1 is accepted. Let $\underline{\theta} = \min\{\theta \in R\}$ where the $\min$ is taken with respect to $\geq_t$.

Then consider the following deviation: take some $t' \neq t$ and let $\alpha(x,t') = \pi_t(x_1|\underline{\theta})$ for all $x \in X$ and $\alpha = \tilde{\alpha}$ otherwise. Because preferences are single-dipped, there is $\theta \in A$ such that $\pi_t(x_1|\theta) < \pi_t(x_1|\underline{\theta})$ and for all $\theta' \in R$, $\pi_t(x_1|\theta') \geq \pi_t(x_1|\underline{\theta})$. Let $s(t'|\theta) = 1$ for that type and $s = \tilde{s}$ otherwise. This deviation is strictly profitable, i.e., $\min_m v(\tilde{\alpha},\tilde{s}) < \min_m v(\alpha,s,m)$.

Suppose the prior is such that $\tilde{\alpha}(x_1,t) = \tilde{\alpha}(x_0,t) \in \{0,1\}$ when only t is offered. This means that the DM does not react to information. Let $\alpha(x,t') = \tilde{\alpha}(x,t)$ for some $t' \neq t$ and

$s(t'|\theta) = 1$ for some $\theta \in A$ and $s = \tilde{s}$ otherwise. We get $\min_m v(\tilde{\alpha}, \tilde{s}) = \min_m v(\alpha, s, m)$, so it is also an solution.

Suppose that the DM's preferences are not single-peaked given $\geq_t$.

In this case, it is possible to find $A_1, A_2 \in A$ and $R_1 \in R$ such that $A_1 <_t R_1 <_t A_2$. Let $\mu(\theta) \approx 0$ for $\theta \neq A_1, A_2, R_1$ and be such that x_0 is rejected and and x_1 is accepted when only t is offered. Because t is informative, there is always such prior. Then from the reasoning above the menu $\{t, t'\}$ is strictly better for the DM than $\{t\}$ when only focusing on A_1, A_2, R_1 have positive probability. But because $\mu(\theta) \approx 0$ for $\theta \neq A_1, A_2, R_1$, then the menu $\{t, t'\}$ remains strictly better than $\{t\}$ whatever the behaviour of the other types.

Suppose that the DM's preferences are not enclosed given $\geq_t$.

If the DM's preferences are not enclosed, then suppose without loss of generality that there is $R_1 \in R$ such that $R_1 \leq_t \theta$ for any $\theta \in \Theta$ (otherwise, simply change the roles of x_1 and x_0).

If for all $\theta \in A$, $\theta =_t R_1$, then preferences are single-peaked and only offering t is optimal.

Suppose it is not the case and take some $A_1, A_2 \in A$ such that $A_1 \geq_t A_1 \geq_t R_1$, with at least one strict inequality. Suppose that for $\theta \neq A_1, A_2, R_1$, $\mu(\theta) \approx 0$. An argument analogue to the proof that single-peakness implies that only t is chosen in equilibrium holds.

A.4 Proof of Proposition 5

Proof. Suppose the DM only uses t and let t' be the coarsened version of t that pools signals in X'. Let $T = \{t, t'\}$. Let $\pi_{t'}(x'|\theta) = \sum_{x \in X'} \pi_t(x|\theta)$ for some $x' \in X'$.

Consider the deviation, $(\tilde{\alpha}, \tilde{s})$: $\tilde{\alpha}(x', t') = \tilde{\alpha}$ and $\tilde{\alpha}(x, \tilde{t}) = \alpha(x, \tilde{t})$ for $x \neq x'$, $\tilde{t} = t, t'$ and $\tilde{s}(t'|\theta) = 1$ if $\sum_{x \in X'} \tilde{\alpha}\pi_t(x|\theta) > \sum_{x \in X'} \alpha(x, t)\pi_t(x|\theta)$and $\tilde{s}(\cdot|\theta) = s(\cdot|\theta)$ otherwise. We want to show that

$$\min_m v(\tilde{\alpha}, \tilde{s}, m) \geq \min_m v(\alpha, s, m)$$
$$\Leftrightarrow \sum_{\theta \in A} \sum_{x \in X'} \mu(\theta) \big[(\tilde{\alpha} - \alpha(x,t))\pi_t(x|\theta)\big]^+ \geq \sum_{\theta' \in R} \sum_{x \in X'} \mu(\theta') \big[(\tilde{\alpha} - \alpha(x,t))\pi_t(x|\theta')\big]^+$$

which is exactly the condition in Proposition 5. Note that the strategy of the R-types is to mimick a type choosing t' iff $\sum_{x \in X'} \tilde{\alpha}\pi_t(x|\theta') > \sum_{x \in X'} \alpha(x,t)\pi_t(x|\theta')$. $\square$

A.5 Proof of Proposition 3

Proof. Note that in an MLRP environment, the strategy of the DM takes the form of a cutoff strategy. For each test t, there is $x_t \in X$ such that $\alpha(x,t) = 0$ for $x < x_t$, $\alpha(x,t) = 1$ for $x > x_t$ and $\alpha(x_t, t) \in [0,1]$. From Lemma 1, we know that there is an optimal menu containing the Blackwell most informative test. Because all tests are MLRP and the DM's payoffs satisfy single-crossing condition, the Lehmann order is well-defined and the Blackwell order implies the Lehmann order (Lehmann, 1988; Persico, 2000). Let $\succeq^a$ denote the Lehmann order.

The Lehmann order is defined on continuous information structure. But as outlined in Lehmann (1988), we can always make our conditional probabilities continuous by adding independent uniform between each signal. Let's assume, without loss of generality, that $X = \{1, ..., n\}$. The new distribution over signal is $\tilde{y}|\theta = \tilde{x}|\theta - u$ where $u \sim$U$[0,1]$. Denote by F_t the cdf associated with the new information structure.

We have that $t \succeq^a t'$ if $y^*(\theta, y) \equiv F_t(y^*|\theta) = F_{t'}(y|\theta)$ is nondecreasing in θ for all y (Lehmann, 1988). In particular, this condition implies that if $F_t(y|\theta') \leq (<) F_{t'}(y'|\theta')$ then $F_t(y|\theta) \leq (<) F_{t'}(y'|\theta)$ for all $\theta > \theta'$.

Let α be the optimal strategy and x_t be the cutoff signal associated to each test. To each

$\big(\alpha(\cdot,t),x_t\big)$ we can associate a $y_t \equiv x_t - \alpha(x_t,t)$.

If t is part of an optimal menu, it must be that there is some $\theta' \in R$ such that $p_t(\alpha;\theta') \geq p_{t'}(\alpha;\theta')$ for all t'. Or put differently, $F_t(y_t|\theta') \leq F_{t'}(y_{t'}|\theta')$ for all t'. But then $F_t(y_t|\theta) \leq F_{t'}(y_{t'}|\theta)$ for all t' and all $\theta > \theta'$, in particular all $\theta \in A$. Therefore all type in A prefer test t as well and there is an solution of the max-min problem where all types in $\theta \in A$ choose t. (If there is an A-type that is indifferent between t and t' then all types in R must be indifferent or prefer t' so choosing t is an equilibrium strategy for such A-type.) □

A.6 Proof of Proposition 6

I first show that if $t > t'$, then $\mu(\cdot|t,x) \succeq_{FOSD} \mu(\cdot|t',x)$ where $\succeq_{FOSD}$ denotes first-order stochastic dominance.

Proof. The proof is similar to the one in Milgrom (1981). Denote by $G_t(\cdot|x)$ the cdf of posterior beliefs after signal x in test t. For all $\theta > \theta'$,

$$\mu(\theta)\frac{\pi_t(x|\theta)}{\pi_t(x|\theta')} \geq \mu(\theta)\frac{\pi_{t'}(x|\theta)}{\pi_{t'}(x|\theta')}$$

Take some $\theta^* \geq \theta'$. Summing over θ, we get

$$\sum_{\theta>\theta^*}\mu(\theta)\frac{\pi_t(x|\theta)}{\pi_t(x|\theta')} \geq \sum_{\theta>\theta^*}\mu(\theta)\frac{\pi_{t'}(x|\theta)}{\pi_{t'}(x|\theta')}$$

Inverting and summing over θ', we get

$$\frac{\sum_{\theta^*\geq\theta'}\mu(\theta')\pi_t(x|\theta')}{\sum_{\theta>\theta^*}\mu(\theta)\pi_t(x|\theta)} \leq \frac{\sum_{\theta^*\geq\theta'}\mu(\theta')\pi_{t'}(x|\theta')}{\sum_{\theta>\theta^*}\mu(\theta)\pi_{t'}(x|\theta)}$$

which implies

$$\frac{G_t(\theta^*|x)}{1-G_t(\theta^*|x)} \leq \frac{G_{t'}(\theta^*|x)}{1-G_{t'}(\theta^*|x)} \quad \Rightarrow \quad G_t(\theta^*|x) \leq G_{t'}(\theta^*|x)$$

$\square$

The way this proof proceeds is by fixing a menu and dividing tests in two categories: (1) those for which $\alpha(x_0,\tilde{t}) \in (0,1)$ and $\alpha(x_1,\tilde{t}) = 1$ and (2) $\alpha(x_0,\tilde{t}) = 0$ and $\alpha(x_1,\tilde{t}) \in (0,1]$. I exclude the possibility that the DM always accepts or rejects after any signal as it would either be the only test chosen in equilibrium or never chosen. Then, I show that within each category, it is without loss of optimality to have at most one test. It is thus optimal to have at most two tests in the menu. The last part of the proof shows that the resulting menu is dominated by having only one test.

If there are two tests, $t > t'$ such that $\alpha(x_0,\tilde{t}) = 0$ and $\alpha(x_1,\tilde{t}) \in (0,1]$, I will show that,

$$p_t(\alpha;\theta') \geq p_{t'}(\alpha;\theta') \quad \Rightarrow \quad p_t(\alpha;\theta) \geq p_{t'}(\alpha;\theta) \text{ for all } \theta > \theta'$$

Take two tests such that $\alpha(x_0,\tilde{t}) = 0$, $t > t'$. Let α, α' denote their respective probability of accepting after x_1. Define $\alpha(\theta) \equiv \alpha(\theta)\pi_t(x_1|\theta) - \alpha'\pi_{t'}(x_1|\theta) = 0$. Rearranging, $\alpha(\theta) = \alpha'\frac{\pi_{t'}(x_1|\theta)}{\pi_t(x_1|\theta)}$. From our assumption on the difficulty environment, $\alpha(\theta)$ is decreasing in θ. If $p_t(\alpha;\theta') \geq p_{t'}(\alpha;\theta')$ for some θ' then $\alpha \geq \alpha(\theta')$. Then $\alpha \geq \alpha(\theta)$ for all $\theta > \theta'$.

In equilibrium, we must have that there is one $\theta' \in R$ that chooses t and thus for all $\theta \in A$, $p_t(\alpha;\theta) \geq p_{t'}(\alpha;\theta)$. Then there is an solution of the max-min problem where t' is never chosen.

A similar argument can be made for all tests where $\alpha(x_0,\tilde{t}) > 0$.

Thus we conclude that it is without loss of optimality that the optimal menu has at most two tests.

Suppose the optimal menu uses two tests, $t > t'$. I will now show that it must be that $\alpha(x_0, t) \in (0, 1)$ and $\alpha(x_1, t') \in (0, 1)$, i.e., the DM must accept in the hard test when there is a fail grade and only accept in the easy test if there is a pass grade. Suppose it is not the case and denote by α, α' their respective mixing probabilities. Define $\alpha(\theta) \equiv \alpha(\theta)\pi_t(x_1|\theta) - \alpha'\pi_{t'}(x_0|\theta) - \pi_{t'}(x_1|\theta) = 0$, which is equivalent to $\alpha(\theta) = \alpha'\frac{1}{\pi_t(x_1|\theta)} + (1-\alpha')\frac{\pi_{t'}(x_1|\theta)}{\pi_t(x_1|\theta)}$. Again from our assumptions, this is decreasing in θ. A type θ chooses t if $\alpha \geq \alpha(\theta)$. Thus if one $\theta \in A$ chooses t all $\theta \in R$ choose t and there is no pooling of A and R-types on t', or it is payoff equivalent to just offering t. Therefore, $\alpha(x_0, t) \in (0, 1)$ and $\alpha(x_1, t') \in (0, 1)$ for $t > t'$.

If the DM mixes, he must be indifferent and thus we have

$$\sum_{\theta \in A} \mu(\theta)\sigma(t|\theta)\pi_t(x_0|\theta) - \sum_{\theta' \in R} \mu(\theta')\sigma(t|\theta')\pi_t(x_0|\theta') = 0$$
$$\sum_{\theta \in A} \mu(\theta)\sigma(t'|\theta)\pi_{t'}(x_1|\theta) - \sum_{\theta' \in R} \mu(\theta')\sigma(t'|\theta')\pi_{t'}(x_1|\theta') = 0$$

In the easy test, because the DM rejects with positive probability after x_1 and rejects for sure after x_0 (as he uses a cutoff strategy), his payoffs from t' is 0, i.e., he does as well as rejecting for sure.

In the hard test, he accepts with some probability after x_0 and thus his payoffs are

$$\sum_{\theta \in A} \mu(\theta)\sigma(t|\theta) - \sum_{\theta' \in R} \mu(\theta')\sigma(t|\theta')$$

that is the payoffs he would get from accepting all types choosing t. Thus the overall payoffs from the menu is $\sum_{\theta \in A} \mu(\theta)\sigma(t|\theta) - \sum_{\theta' \in R} \mu(\theta')\sigma(t|\theta')$. Offering a menu is better than a singleton menu if this value is strictly greater than offering t and following the signal

$$\sum_{\theta\in A}\mu(\theta)\sigma(t|\theta)-\sum_{\theta'\in R}\mu(\theta')\sigma(t|\theta') > \sum_{\theta\in A}\mu(\theta)\pi_t(x_1|\theta)-\sum_{\theta'\in R}\mu(\theta')\pi_t(x_1|\theta')$$
$$= \sum_{\theta\in A}\sigma(t|\theta)\mu(\theta)\pi_t(x_1|\theta)+\sum_{\theta\in A}\sigma(t'|\theta)\mu(\theta)\pi_t(x_1|\theta)$$
$$-\sum_{\theta'\in R}\sigma(t|\theta')\mu(\theta')\pi_t(x_1|\theta')-\sum_{\theta'\in R}\sigma(t'|\theta')\mu(\theta)\pi_t(x_1|\theta')$$

We can rearrange and use the indifference condition at (x_0, t) to get

$$0 > \sum_{\theta\in A}\sigma(t'|\theta)\mu(\theta)\pi_t(x_1|\theta)-\sum_{\theta'\in R}\sigma(t'|\theta')\mu(\theta)\pi_t(x_1|\theta')$$

Using the indifference condition at (x_1, t'), we can replace 0 on the LHS and get

$$\sum_{\theta\in A}\mu(\theta)\sigma(t'|\theta)\pi_{t'}(x_1|\theta)-\sum_{\theta'\in R}\mu(\theta')\sigma(t'|\theta')\pi_{t'}(x_1|\theta')$$
$$> \sum_{\theta\in A}\sigma(t'|\theta)\mu(\theta)\pi_t(x_1|\theta)-\sum_{\theta'\in R}\sigma(t'|\theta')\mu(\theta)\pi_t(x_1|\theta')$$

But from the definition of the environment, for all $\theta > \theta'$,

$$\frac{\pi_t(x_1|\theta)}{\pi_t(x_1|\theta')} \geq \frac{\pi_{t'}(x_1|\theta)}{\pi_{t'}(x_1|\theta')}$$

which implies that $\mu(\theta|x_1, t) \succeq_{FOSD} \mu(\theta|x_1, t')$. Thus we get a contradiction.

A.7 Proof of Proposition 7

Suppose condition (2) holds. Suppose $(\alpha, s) \in \arg\max\min_m v(\alpha, s, m)$ and $s(t_j|\theta) = 1$ for all $\theta \in A$. Take $(\tilde{\theta}_i, \tilde{\theta}_j) \in \arg\min_{\theta\in A} p_{t_j}(\alpha;\theta)$. Because $p_{t_j}(\alpha;\theta_i,\theta_j)$ is constant in θ_i, we have $(\overline{\theta}_i, \tilde{\theta}_j) \in \arg\min_{\theta \in \Theta} p_{t_j}(\alpha; \theta)$ as well and from condition (2), $(\overline{\theta}_i, \tilde{\theta}_j) \in A$. Consider the deviation to $(\tilde{\alpha}, \tilde{s})$ such that for t_i,

- $\tilde{\alpha}(\cdot, t_i)$ is set so that it has a cutoff structure and $p_{t_i}(\tilde{\alpha}|\overline{\theta}_i, \tilde{\theta}_j) = p_{t_j}(\alpha; \overline{\theta}_i, \tilde{\theta}_j)$ and $\tilde{\alpha}(\cdot, t_j) = \alpha(\cdot, t_j)$ otherwise.
- $\tilde{s}(t_i|\overline{\theta}_i, \tilde{\theta}_j) = 1$ and $\tilde{s}(\cdot|\theta) = s(\cdot|\theta)$ otherwise.

Because the test t_i has the strict MLRP when restricting attention to dimension i, for all $\theta_i < \overline{\theta}_i$, $\min_{\theta \in \Theta} p_{t_j}(\alpha; \theta) \geq p_{t_i}(\tilde{\alpha}|\overline{\theta}_i, \theta_j) > p_{t_i}(\tilde{\alpha}|\theta_i, \theta_j)$. This means that mimicking $(\overline{\theta}_i, \tilde{\theta}_j)$ is weakly dominated and $(\overline{\theta}_i, \tilde{\theta}_j)$ has the probability of being accepted. Thus $\min_m v(\alpha, s, m) \leq \min_m v(\tilde{\alpha}, \tilde{s}, m)$.

Suppose condition (2) does not hold.

If condition (2) is not satisfied, then there a dimension, say 1, and $\tilde{\theta}_2 \in \Theta_2$ such that $(\overline{\theta}_1, \tilde{\theta}_2) \in R$. By the definition of the bidimensional environment, this implies that $(\theta_1, \tilde{\theta}_2) \in R$ for all $\theta_1 \in \Theta_1$. Moreover, for all $\theta_2 < \tilde{\theta}_2$ and all $\theta_1 \in \Theta_1$, $(\theta_1, \theta_2) \in R$.

Now suppose μ is such that $\mu(\theta_1, \tilde{\theta}_2) > \sum_{\theta_2' \neq \theta_2} \mu(\theta_1, \theta_2')$ for all $\theta_1 \in \Theta_1$. And that $\mu(\theta_1, \theta_2) \approx 0$ for all $(\theta_1, \theta_2) \in R$ such that $\theta_2 > \tilde{\theta}_2$.

I am going to show that $\{t_2\}$ is optimal when t_1 fully reveals dimension 1. Because this test can replicate the strategies of any t_1, it is enough to prove our claim.

Suppose there is an optimal menu $\{t_1, t_2\}$. From our assumptions on μ, the DM follows a cutoff strategy after t_2. That's because his payoff is monotone along that dimension, ignoring $(\theta_1, \theta_2) \in R$ such that $\theta_2 > \tilde{\theta}_2$ whose prior probability is close to zero. So it does not upset the cutoff structure of the best-response. This implies that $p_{t_2}(\alpha; \theta_1, \theta_2) > p_{t_2}(\alpha; \theta_1, \tilde{\theta}_2)$ for all $\theta_2 > \tilde{\theta}_2$ because the likelihood ratio is strictly increasing.

Suppose that some $(\theta_1, \tilde{\theta}_2)$ chooses t_1 with probability 1 in equilibrium. Because $\mu(\theta_1, \tilde{\theta}_2) > \sum_{\theta_2' \neq \theta_2} \mu(\theta_1, \theta_2')$ for all $\theta_1 \in \Theta_1$, it must be that the best-response is $\alpha(x = \theta_1, t_1) = 0$ (recall that t_1 fully reveals θ_1). Thus $p_{t_2}(\alpha; \theta_1, \theta_2) = 0$ for all $\theta_2 \in \Theta_2$, otherwise there is a profitable deviation. Either this contradicts the fact that the DM best replies or in equilibrium the DM rejects after all signals in every test. But then he is weakly better off only offering t_2.

Thus to have $\{t_1, t_2\}$ strictly better, it must be that all $(\theta_1, \tilde{\theta}_2)$ choosing t_1 mix in equilibrium. This means that $p_{t_1}(\alpha; \theta_1, \tilde{\theta}_2) = p_{t_1}(\alpha; \theta_1, \tilde{\theta}_2)$. But by the cutoff structure of $\alpha(\cdot, t_2)$, we have $p_{t_2}(\alpha; \theta_1, \theta_2) \geq p_{t_2}(\alpha; \theta_1, \tilde{\theta}_2)$ for all $\theta_2 > \tilde{\theta}_2$ and $p_{t_2}(\alpha; \theta_1, \theta_2) \leq p_{t_2}(\alpha; \theta_1, \tilde{\theta}_2)$ for all $\theta_2 < \tilde{\theta}_2$. Thus t_1 is weakly dominated in the auxiliary max-min problem for all $(\theta_1, \theta_2) \in A$. Thus choosing only $\{t_2\}$ is an optimal menu.

A.8 Proof of Proposition 8

Proof. I will first prove the following lemma. This result already exists in the literature and I provide a proof for completeness.

Lemma 4. *For any $t \succeq t'$ and $\alpha(\cdot, t')$, there is $\alpha(\cdot, t)$ such that*

$$\sum_x \alpha(x, t)\pi_t(x|\theta) \geq \sum_x \alpha(x, t')\pi_{t'}(x|\theta)$$

$$\text{for all } \theta' \in R, \quad \sum_x \alpha(x, t)\pi_t(x|\theta') \leq \sum_x \alpha(x, t')\pi_{t'}(x|\theta')$$

Proof. We can prove this lemma by using a theorem of the alternative (see e.g., Rockafellar (2015) Section 22). Only one of the following statement is true:

- There exists $\alpha(\cdot,t)$ such that

$$\begin{aligned}
&\sum_x \alpha(x,t)\pi_t(x|\theta) \geq \sum_x \alpha(x,t')\pi_{t'}(x|\theta) \\
\text{for all } \theta' \in R, \quad &\sum_x \alpha(x,t)\pi_t(x|\theta') \leq \sum_x \alpha(x,t')\pi_{t'}(x|\theta') \\
\text{for all } x \in X, \quad &\alpha(x,t) \leq 1 \\
\text{for all } x \in X, \quad &\alpha(x,t) \geq 0
\end{aligned}$$

- There exists $z, y \geq 0$ such that

$$\text{for all } x \in X, \quad -z_\theta \pi_t(x|\theta) + \sum_{\theta' \in R} z_{\theta'} \pi_t(x|\theta') + y_x \geq 0 \tag{5}$$

$$-z_\theta \sum_{x'} \alpha(x',t')\pi_{t'}(x'|\theta) + \sum_{\theta' \in R} z_{\theta'} \sum_{x'} \alpha(x',t')\pi_{t'}(x'|\theta') + \sum_{x'} y_{x'} < 0 \tag{6}$$

Take inequality (5) from the second alternative and multiply by $\beta(x,x')$ as described in Definition 6 and sum over $x \in X$:

$$-z_\theta \sum_x \beta(x,x')\pi_t(x|\theta) + \sum_{\theta' \in R} z_{\theta'} \sum_x \beta(x,x')\pi_t(x|\theta') + \sum_x \beta(x,x') y_x \geq 0$$

Because $t \succeq_\theta t'$, we get for all $x' \in X$,

$$-z_\theta \pi_{t'}(x'|\theta) + \sum_{\theta' \in R} z_{\theta'} \pi_{t'}(x'|\theta') + \sum_x \beta(x,x') y_x \geq 0$$

We can then multiply by $\alpha(x', t')$ and sum over $x' \in X$:

$$-z_\theta \sum_{x'} \alpha(x', t')\pi_t(x'|\theta) + \sum_{\theta' \in R} z_{\theta'} \sum_{x'} \alpha(x', t')\pi_{t'}(x'|\theta') + \sum_{x,x'} \alpha(x', t')\beta(x, x')y_x \geq 0 \quad (7)$$

Because $\sum_{x'} \beta(x, x') \leq 1$ and $\alpha(x', t') \leq 1$ for all $x' \in X$, we have $\sum_{x,x'} \alpha(x', t')\beta(x, x')y_x \leq \sum_x y_x$. Therefore, the inequality (6) cannot hold and the first alternative holds. □

With this result in hand, we can now prove our result. Suppose that t is not part of the optimal menu. Thus we can find an solution of the max-min problem of Theorem 1, (α, s) with $s(t|\theta) = 0$ for all $\theta \in A$. Take a test t' used in the solution by some $\theta \in A$. Then from Lemma 4, we can construct a $\tilde{\alpha}$ such that

$$\begin{aligned} & p_t(\tilde{\alpha}; \theta) \geq p_{t'}(\alpha; \theta) \\ \text{for all } \theta' \in R, \quad & p_t(\tilde{\alpha}; \theta') \leq p_{t'}(\alpha; \theta') \end{aligned}$$

If the first inequality is strict or the second such that $m(\theta'|\theta) > 0$ at some minimiser of v at $\tilde{\alpha}$ is strict then we have a strict profitable deviation. Otherwise, we have constructed a new solution of the max-min problem.

□

A.9 Proof of Proposition 9

Proof. $(\Leftarrow)$ For each $\theta \in A$, let t_θ such that

$$\operatorname{supp} \pi_t(\cdot|\theta) \cap \left(\cup_{\theta' \in R} \operatorname{supp} \pi_t(\cdot|\theta') \right) = \emptyset$$

Then posting a menu $(t_\theta)_{\theta\in A}$ is optimal (eliminating duplicates if there are some). Each $\theta \in A$ chooses t_θ. For any strategy of $\theta' \in R$, the DM accepts after any $(x, t) \in \cup_{\theta:\sigma(t|\theta)=1} \operatorname{supp} \pi_t(\cdot|\theta)$ and rejects otherwise. This gives the DM and the A-types maximal payoffs and the R-types get rejected for any strategy they follow.

($\Rightarrow$) Suppose the DM's payoffs are maximal and there is $\theta \in A$ and for all $t \in T$ there is $\theta' \in R$ and $x \in X$ such that $\pi_t(x|\theta), \pi_t(x|\theta') > 0$. Then when θ chooses t out of the menu of tests, if θ' chooses t as well, at x, either the DM accepts θ' or rejects θ. Therefore, payoffs cannot be maximal. □

A.10 Proof of Theorem 2

The only thing we need prove is that it is optimal to have a different message for each type $\theta \in A$, the rest follows from Theorem 1. Suppose it is not the case and take a solution (α, s) of the max-min problem.

There is $\theta_1, \theta_2 \in A$ and $(t, c) \in T \times C$ such that $s(t, c|\theta_1) = s(t, c|\theta_2) = 1$ (if they use a different test then we can also change the message and nothing is changed). Then consider the alternative strategy α' where, for some unused (t, c') in the original mechanism, $\alpha'(t, c', x) = \alpha(t, c, x)$ for all $x \in X$ and $\alpha'(t'', c'', x) = \alpha'(t'', c'', x)$ for all other $(t'', c'') \in T \times C$ and all $x \in X$ otherwise. The new strategy α' is thus the same as α but makes sure that if the pair (t, c') is chosen, it uses the same actions as (t, c). Now consider the following strategy $\tilde{s}(\cdot|\theta)$ for $\theta \in A$ in the auxiliary max-min problem, $\tilde{s}(\cdot|\theta) = s(\cdot|\theta)$ for $\theta \neq \theta_1$ and $\tilde{s}(t, c'|\theta_1) = 1$. In the max-min problem under the strategy α', the payoffs are the same than under (α, s), fixing m for all types. Moreover, any deviations under α' gives the same payoff than under α. Therefore, $(\alpha', \tilde{s})$ is an solution of the max-min problem.

A.11 Proof of Proposition 10

The way this proof proceed is by first arguing that an optimal mechanism $\tilde{\sigma} : \Theta \to \Delta(T \times C)$ does weakly better than an optimal GR-mechanism, τ. Then I will show that the outcome of the optimal mechanism $\tilde{\sigma}$ can be implemented by a GR-game.

To see the first part, note that a GR-mechanism can be rewritten as a mechanism $\tilde{\tau} : C \to \Delta(T)$ and a DM-strategy $\tilde{\alpha} : C \times T \times X \to [0, 1]$. Then we can implement any equilibrium outcome of $(\tilde{\tau}, \tilde{\alpha}, \delta)$, where δ is the agent's strategy by a mechanism and strategy of the DM, $(\tilde{\sigma}, \alpha)$ by setting $\tilde{\sigma} = \tilde{\tau} \circ \delta$, the composition of the GR-mechanism and the agent's strategy and $\alpha = \tilde{\alpha}$. This does not change the outcome so all the agent's incentives are preserved.

I will now show that the outcome of the menu game with communication can be implemented in a GR-game.

Remember that we have established that in the max-min problem, all the A-types play a pure strategy and send a different message (Theorem 2). This implies that it is without loss of optimality to decompose the A-types' strategy s in choosing a message $c \in C$, call it $\phi : A \to C$ and a test for each message, call it $\rho : C \to T$.

Abusing notation define

$$
\begin{aligned}
p_t(\alpha; \theta, c) &= \sum_x \alpha(t, x, c)\pi_t(x|\theta) \\
v(\alpha, \phi, \rho, m) &= \sum_c \mathbb{1}\big[(t, \theta) : t = \rho(c), c = \phi(\theta)\big] \big[\mu(\theta)p_t(\alpha; \theta, c) \\
&\qquad\qquad\qquad\qquad - \sum_{\theta' \in R} \mu(\theta')m(\theta|\theta')p_t(\alpha; \theta', c)\big]
\end{aligned}
$$

To understand the new version of v, we sum over all messages and for each message, we select the the test associated with it and the A-type choosing that message.

We get,

$$\min_{m} \max_{\alpha,s} v(\alpha,s,m) = \max_{\alpha,s} \min_{m} v(\alpha,s,m) = \max_{\alpha,\phi,\rho} \min_{m} v(\alpha,\phi,\rho,m)$$

But now observe that we could equivalently interpret ρ as being chosen by the DM as it maximises his payoffs. By the same arguments as in Theorem 1, the strategies generate an equilibrium strategy.