Robust decision-making*

Vinicius Carrasco† Humberto Moreira‡

July 10, 2013

Abstract

We consider a setting in which a decision-maker (principal) has to rely on an informed (but biased) agent to make a decision and is uncertain about the distribution from which the relevant state is drawn. The decision-maker is uncertainty averse, i.e., she has a maximim objective. The optimal (robust) mechanism is fully characterized by the property that it ensures the same payoff for the decision-maker across all states with positive likelihood. The shape of the robust mechanism depends on the agent's bias. When the agent's bias is state-dependent, the robust mechanism is typically stochastic and can be interpreted as representing a contingent decision plan. However, if the agent's bias is constant and the set state space is unbounded, the robust mechanism is deterministic and entails full delegation.

Keywords: decision-making, robustness, asymmetric information, delegation.

J.E.L. Classifications: D82 (Asymmetric and Private Information), D86 (Economics of Contract Theory).

1 Introduction

A CEO relies on report by a division manager before deciding whether to invest in a given project; a regulator sets prices for a regulated firm according to information it provides about its costs; and an Economics Department decides which candidates to fly out for job talks based on information provided by a recruiting committee. These are only a few of many practical examples of situations in which a decision-maker (DM or principal) has to rely on a privately informed and potentially biased agent to make a decision.

Being applicable to a wide range of economic situations, it is then no surprise that a large literature in Economics has studied models of decision-making under asymmetric information. A common feature of virtually all papers on decision-making under asymmetric information is the assumption that, despite being uninformed about an underlying ex-post state (observed by the agent), the DM fully knows the distribution from which such state is drawn.[1] Yet, there are relevant instances in which the DM might not know the

*A preliminary version of this paper was presented at Universidad de Chile (October, 2010). We have benefited from comments by Felipe Balmaceda, Juan Escobar and Nicolas Figueroa. We also thank Andy Skrzypacz for useful conversations. All errors are ours.

†Department of Economics, PUC-Rio. Rua Marquês de São Vicente, 225, Gávea. Rio de Janeiro, Brazil. Email address: vnc@econ.puc-rio.br

‡FGV\EPGE Escola Brasileira de Economia e Finanças. Praia de Botafogo, 190, 22250-900. Rio de Janeiro, Brazil. Email address: humberto@fgv.br

[1] Alternatively, the DM is able to form a unique prior regarding those states.

distribution of states. These situations are likely to be the cases for once in a lifetime decisions, or when a decision is to be made for the first time, or under few observations of previous decisions made. In those cases, what decision-making procedure should be implemented?

To a standard model in which a committed principal interacts with a biased agent who is privately informed about a relevant state, this paper adds two features: (i) the principal does not know the distribution from which the relevant state is drawn; and (ii) in face of this uncertainty, the principal is uncertainty averse, i.e., she[2] has a maximin objective. This last assumption corresponds to considering a DM who, being ignorant about the likelihood of different states, wishes to design a decision-making procedure that does well in a wide range of circumstances.

As the majority of papers in the literature, we derive our results for a setting in which both the agent's and the principal's ex-post payoffs are quadratic. Hence, whenever confronted with general stochastic mechanisms, they only care about two features of these mechanisms: expected decision and variance. As shown by Kováč and Mylanov (2009), one can use standard envelope arguments to characterize the set of incentive compatible stochastic mechanisms. We fully describe the principal's ex-post payoff in an incentive compatible mechanism in terms of the expected decision implied by the mechanism and the agent's payoff at the worst possible state. This significantly simplifies the task of finding the robust decision-making mechanism. The second step of our derivation makes use of the following equivalence result. Finding a robust decision-making mechanism is equivalent to computing the set of Nash equilibria of a simultaneous zero-sum game played by the principal – whose choice set is the set of incentive compatible stochastic mechanisms – and the adversary nature, who picks the distribution of states to minimize the principal's expected payoff.[3]

Since nature's objective is to minimize the principal's (expected) payoff, in equilibrium, it will only place positive weight on states for which the principal's payoff is the lowest possible value. Therefore, the principal's payoff will be constant over all states with positive likelihood. While the game theoretical interpretation of these results is standard (the principal makes his choice to make nature indifferent between any "action" in the support of its strategy), the interpretation in terms of the robust mechanism is quite new and interesting. The principal insures against uncertainty by designing a mechanism in which her payoff does not depend on the realized state. The requirement that the principal's payoff remains constant across all states with positive likelihood leads to an ordinary differential equation (ODE) that fully pins down the behavior of the expected decision in a robust scheme. Much as in mechanism design problems with transferable utility, one can then recoup the variance of the mechanism from the expected decision using the usual envelope condition that, along with the monotonicity condition, characterizes incentive compatibility. Given that the principal's equilibrium behavior is pinned down by nature's "indifference" condition, nature selects the distribution of states to guarantee that the principal is willing to behave according to what described above. This, in turn, also leads to another ODE that the "worst-case" distribution has to satisfy. Particular solutions of both ODEs fully characterize a Nash equilibrium of the zero-sum game and, hence, a robust mechanism. Particular features of the robust scheme are shown to depend on the form of the agent's bias, which is the difference between her most preferred decision and the principal's.

When the bias is constant and the state space is unbounded, the robust mechanism is deterministic and the recommended decision coincides with the agent's most preferred decision in every state. Since the robust mechanism is deterministic, it can be (indirectly) implemented through a delegation set and, since the

[2] We will use feminine pronouns to refer to principal and masculine pronouns in referring to the agent.

[3] In the zero-sum game, both the principal and the nature are von-Neumann and Morgenstern expected utility maximizers.

decision made always coincides with the agent's most preferred one, the (indirect) implementation entails full delegation. Therefore, the principal grants full discretion to the agent to make the decision.

When the bias varies with the state, the robust mechanism is stochastic and cannot be indirectly implemented through a delegation set. One may then interpret the decision process as being contingent on the realization of a (payoff irrelevant) random device that, somewhat surprisingly, plays the role of guaranteeing that the principal is fully insured against uncertainty. Put differently, the best way the principal finds to insure against the opportunistic behavior of the adversary nature is by forcing nature to face itself some "uncertainty", implied by a random decision-making process. While without further specification of the bias function we cannot fully characterize the robust decision-making mechanism, one can quite generally show that the expected decision is strictly between the principal's and the agent's favorite ones. Also, we can show that the degree of the responsiveness of the expected decision to the state will be larger (i) the larger the responsiveness of principal's favorite decision to the state, and (ii) the smaller the agent's bias.

Related literature

We now discuss how our findings compare to the existing literature. As mentioned earlier, there is a large literature on (Bayesian) decision-making under asymmetric information for settings similar to ours, in which a principal can commit to mechanisms and side payments are not allowed.[4] Holmström (1977, 1984), Melumad and Shibano (1991), Alonso and Matoushek (2008), and Kováč and Mylanov (2009) all fit this description. In contrast to our work, all those papers assume that the principal is an expected utility maximizer who knows (or, alternatively, holds a single prior of) the distribution from which states are drawn. Our alternative behavioral assumption leads to differences in the results we obtain. First, the fact that the principal designs a mechanism to guarantee the same payoff for the DM across all states that might occur is a unique implication of a maximin behavior. In fact, as we have argued, this is the form by which the principal insures against uncertainty. Second, full delegation is not, in general, optimal in those papers. In fact, to the best of our knowledge, only Alonso and Matoushek (2008) have a full delegation result for the "normal-linear" version of their model when the agent's bias is sufficiently small (i.e., when the agent is "moderate" in their jargon).[5] Third, as shown by Kováč and Mylanov (2009), under mild regularity conditions, stochastic mechanisms are not optimal in Bayesian decision-making problems when both the agent's and the principal's preferences are quadratic.[6]

There is a growing literature in robust mechanism design. Ely and Chung (2007), for example, show that a selling mechanism robust to buyers' beliefs must be implementable in dominant strategies. Bergmann and Shlag (2009) consider the design of selling mechanisms that minimize seller's maximum regret. Analogously to the optimality of stochastic mechanisms for the case of non-constant bias in our model, they show that the optimal scheme leads to a stochastic pricing mechanism. Hurwicz and Shapiro (1978) establish that

[4]Starting with Crawford and Sobel (1981), there is also an extensive literature dealing with the case in which the principal lacks commitment. The main interest of this literature is to understand the amount of communication that can take place between the agent and the principal for varying communication protocols (e.g., one round communication vis à vis multiple rounds).

[5]Dessein (2002) considers a principal's choice between two given mechanisms – full delegation or the one-round communication protocol of Crawford and Sobel (1981) –, and shows that full delegation does better when the agent's bias is small.

[6]Kováč and Mylanov (2009) prove that stochastic mechanisms are optimal whenever the principal's payoff is the absolute value of the difference between her most favorite decision and the decision made. For that case, whenever the sign of the difference between her most favorite decision and the decision made does not change, the principal is de facto risk neutral.

an equal division of the crop may be the robust sharecropping mechanism, whereas Carroll (2013) shows that linear incentive contracts are robust to a principal's uncertainty regarding the technology in a moral hazard setting. In this literature, the only paper dealing with decision-making problems is Frankel (2013). He considers the problem faced by a maximin principal in designing multidimensional decision rules when, although knowing the distribution of states, she is uncertain about the agent's *preferences.* Frankel shows quite generally that implementation through delegation sets is optimal and, in particular, robust mechanisms are deterministic. The sources of uncertainty this paper deal with are quite different from ours. While he focuses on the uncertainty regarding agent's preferences, in our one-dimensional problem, we assume that the principal knows agent's preferences for sure but is uncertain about the distribution of states. It is then no surprise that the results also differ. First, whenever the robust mechanism is stochastic in our setting, implementation cannot be done through delegation sets. Also, full delegation is never optimal in his setting, while it is in ours for the constant bias case and unbounded state space.

2 Model

There are two players, an agent and a principal. The principal must make a decision a chosen from the set $[\underline{a}, \overline{a}] := \mathcal{A} \subset \mathbb{R}$. Both the principal and the agent's preferences depend on the decision made and an underlying state $\theta \in [0,1]$, which is privately observed by the agent. Although (as we describe below) ignorant about the distribution from which states are drawn, the principal knows that all possible states are in $[0,1]$.

The principal's and agent's ex-post (i.e., after the state θ realizes and the decision a is made) utilities are

$$u_p(a,\theta) = -\left(a - y_p\left(\theta\right)\right)^2 \text{ and } u_a(a,\theta) = -\left(a - y_a\left(\theta\right)\right)^2,$$

where $y_p\left(\theta\right)$ and $y_a\left(\theta\right)$ are strictly increasing and continuously differentiable functions.

From their preferences, given a state θ, the principal's most preferred decision is $y_p\left(\theta\right)$, whereas the agent's is $y_a\left(\theta\right)$. We assume that the set $\mathcal{A}$ is large enough to contain both players' favorite decisions for all states $\theta \in [0,1]$. We define the difference of their favorite actions as the agent's bias:

$$b\left(\theta\right) = y_a\left(\theta\right) - y_p\left(\theta\right).$$

Throughout, we assume that $b\left(\theta\right)$ is strictly positive and non-decreasing. Hence, the agent systematically prefers higher decisions than the principal and their misalignment of interests does not decrease with the state.

Assumption A. *The bias function $b\left(\theta\right)$ is strictly positive, continuous and non-decreasing.*

In contrast to all the literature in decision-making with asymmetric information, we assume that, on top of being uninformed about the state, the principal does not know with certainty the distribution from which the state θ is drawn (or, alternatively, she does not hold a unique prior regarding the state θ). In face of this uncertainty, the principal has a "maximin" objective, that is, she maximizes her expected utility for the worst among all distributions.

2.1 Mechanisms

The principal can commit to any mechanism she offers to the agent. By the Revelation Principle, she can restrict attention to incentive compatible direct mechanisms. It is convenient to allow for stochastic mechanisms, that is, mechanisms taking announcements made by the agent on probability distributions over decisions. Formally, letting $\mathbb{P}(\mathcal{A})$ be the set of all probability measures on $\mathcal{A}$, a direct stochastic mechanism is a collection $(dP_\theta(\cdot))_{\theta\in[0,1]} \subset \mathbb{P}(\mathcal{A})$ of probability distributions over decisions.

Given a direct stochastic mechanism $(dP_\theta(\cdot))_{\theta\in[0,1]}$, the agent's payoff in case he announces $\hat{\theta}$ and the state is θ is

$$\int_{\mathcal{A}} u_a(a,\theta)dP_{\hat{\theta}}(a), \tag{1}$$

whereas the principal's is

$$\int_{\mathcal{A}} u_p(a,\theta)dP_{\hat{\theta}}(a). \tag{2}$$

Define

$$\overline{a}(\hat{\theta}) = \int_{\mathcal{A}} a dP_{\hat{\theta}}(a) \tag{ED}$$

as the *expected decision* and

$$\sigma^2(\hat{\theta}) = \int_{\mathcal{A}} (a - \overline{a}(\hat{\theta}))^2 dP_{\hat{\theta}}(a) \tag{V}$$

as the *variance* of the decision implied by the mechanism. Therefore, (1) can be rewritten as

$$U_a^{(\overline{a},\sigma^2)}(\hat{\theta}|\theta) = \int_{\mathcal{A}} u_a(a,\theta)dP_{\hat{\theta}}(a) = u_a(\overline{a}(\hat{\theta}),\theta) - \sigma^2(\hat{\theta})$$

and, since $u_p(a,\theta) = u_a(a,\theta) - 2b(\theta)(a - y_a(\theta)) - b(\theta)^2$ for all a and θ, (2) as

$$\begin{aligned} U_p^{(\overline{a},\sigma^2)}\left(\hat{\theta}|\theta\right) &= \int_{\mathcal{A}} [u_a(a,\theta) - 2b(\theta)(a - y_a(\theta)) - b(\theta)^2] dP_{\hat{\theta}}(a) \\ &= U_a^{(\overline{a},\sigma^2)}(\hat{\theta}|\theta) - 2b(\theta)(\overline{a}(\hat{\theta}) - y_a(\theta)) - b(\theta)^2 . \end{aligned} \tag{3}$$

Abusing the notation, we also denote the agent's and principal's payoffs when the agent reveals the true state by

$$U_a^{(\overline{a},\sigma^2)}(\theta) = U_a^{(\overline{a},\sigma^2)}(\theta|\theta) \text{ and } U_p^{(\overline{a},\sigma^2)}(\theta) = U_p^{(\overline{a},\sigma^2)}(\theta|\theta).$$

As noticed by Kováč and Mylanov (2009), with quadratic preferences the variance of the decision-making process (henceforth, the "variance" of the mechanism) and the expected decision implied by $(dP_\theta(\cdot))_{\theta\in[0,1]}$ matter for both the principal and the agent. We can then reduce any general direct stochastic mechanism by its first and second moments. For tractability, we restrict our attention to the set of direct stochastic mechanisms such that their first and second moments are integrable functions of the state. This observations motivate the following definition.

Definition 1. *The space of direct mechanisms is given by*

$$\mathcal{M} = \{m := (\overline{a},\sigma^2) : [0,1] \to \mathcal{A} \times \mathbb{R}_+ \text{ integrable function}\}.$$

Principal's problem

The principal's problem is one of choosing a direct mechanism in $\mathcal{M}$ to maximize her expected utility under the worst-case distribution of states subject to incentive compatibility constraints. Letting $\mathcal{F}$ be the set of all distributions over $[0,1]$, one can write her problem as

$$\begin{aligned} &\max_{m=(\overline{a},\sigma^2)\in\mathcal{M}} \min_{F\in\mathcal{F}} \left[\int_0^1 U_p^m(\theta)\, dF(\theta)\right] \\ &s.t.\ U_a^m(\theta) \geq u_a(\overline{a}(\hat{\theta}),\theta) - \sigma^2(\hat{\theta}),\ \text{ for all } \theta,\hat{\theta}\in[0,1]\,. \end{aligned} \tag{4}$$

Definition 2. *A solution of problem (4) is called a robust decision-making mechanism.*

When a privately informed agent has quadratic preferences in a decision-making problem with non-transferable utility, the variance of the mechanism, $\sigma^2(\theta)$, plays an analogous role in terms of slacking incentive compatibility constraints as transfers in a setting where side payments are available and preferences are quasi-linear. The only difference one must take into account is the fact that the variance has to be non-negative. Once this latter constraint is considered, the following characterization of incentive compatibility, which also appears in Kováč and Mylanov (2009), ensues:

Lemma 1 *A mechanism $m = (\overline{a}, \sigma^2) \in \mathcal{M}$ is incentive compatible if and only if, for all $\theta \in [0,1]$,*

(i) $U_a^m(\theta) = U_a^m(0) + 2\int_0^\theta (\overline{a}(\widetilde{\theta}) - y_a(\widetilde{\theta}))y_a'(\widetilde{\theta})d\widetilde{\theta}$;

(ii) $\overline{a}(\theta)$ *is non-decreasing;*

(iii) $\sigma^2(\theta) = u_a(\overline{a}(\theta),\theta) - U_a^m(\theta) \geq 0.$

If $\overline{a}(\theta) \leq y_a(\theta)$, *for all* $\theta\in[0,1]$, *then* $\sigma^2(\theta)$ *is a non-decreasing, and (iii) is equivalent to*

(iii') $\sigma^2(0) = u_a(\overline{a}(0),0) - U_a^m(0) \geq 0.$

Lemma 1 is standard in the mechanism design literature. Since the agent's preferences satisfy a single-crossing condition in $(\overline{a},\theta)$, one can replace the incentive compatibility constraints by the integral representation of the agent's utility implied by the Envelope Theorem (Milgrom and Segal, 2002) and a monotonicity condition that the expected decision has to satisfy. In addition to those two standard conditions, one has to take into consideration a non-negativity constraint on the variance of the mechanism.

The last piece of Lemma 1 says that, for any mechanism for which the expected decision lies weakly below the agent's most preferred decision, the non-negativity constraint on the variance will be satisfied if and only if it holds at the lowest possible state 0. This last condition is useful for what follows because, as we establish in subsection 3.3, the robust decision-making mechanism satisfies this property.

3 Robust decision-making mechanism

To characterize the robust decision-making mechanism, it is useful to think of problem (4) as the following game. There are two players. The principal whose (von-Neumann and Morgenstern) payoff is

$$u_p(a,\theta) = -(a - y_p(\theta))^2$$

and chooses a mechanism $m \in \mathcal{M}$, and the nature (n) whose (von-Neumann and Morgenstern) payoff is

$$u_n(a, \theta) = (a - y_p(\theta))^2 = -u_p(a, \theta),$$

and chooses a distribution $F \in \mathcal{F}$. Clearly, this is a zero-sum game as $u_p(a, \theta) + u_n(a, \theta) = 0$, for all $a \in \mathcal{A}$ and $\theta \in [0, 1]$. The following result is a widely known property of zero-sum games (see, for example, von-Neumann and Morgenstern (1944) and Osborne and Rubinstein (1994)).

Lemma 2 *If $(F^*, m^*) \in \mathcal{F} \times \mathcal{M}$ is a Nash equilibrium of the game the principal plays against nature, then $m^* \in \mathcal{M}$ is a robust decision-making mechanism. Moreover,*

$$\max_{m\in\mathcal{M}} \min_{F\in\mathcal{F}} \int U_p^m(\theta) dF(\theta) = \int U_p^{m^*}(\theta) dF^*(\theta) = \min_{F\in\mathcal{F}} \max_{m\in\mathcal{M}} \int U_p^m(\theta) dF(\theta).$$

From Lemma 2, by establishing that an equilibrium of the game the principal plays against nature exists, one can show that a robust decision-making mechanism exists.

Define the nature's payoff functional in the game when the principal behaves optimally by

$$\Omega(F) = \sup_{m\in\mathcal{M}} \int U_p^m(\theta) dF(\theta),$$

for all $F \in \mathcal{F}$. If $(F^*, m^*) \in \mathcal{F} \times \mathcal{M}$ is a Nash equilibrium of the zero-sum game, then

$$F^* \in \arg\min_{F\in\mathcal{F}} \Omega(F) \text{ and } \Omega(F^*) = \int U_p^{m^*}(\theta) dF^*(\theta).$$

We have the following useful characterization.

Lemma 3 *Let $m^* \in \mathcal{M}$ such that*

$$\int U_p^{m^*}(\theta) dF(\theta) \geq \int U_p^{m^*}(\theta) dF^*(\theta) = \Omega(F^*), \text{ for all } F \in \mathcal{F}. \tag{5}$$

Then, F^* minimizes Ω on $\mathcal{F}$ and $m^* \in \mathcal{M}$ is a robust decision-making mechanism.

Proof. By the definition of Ω

$$\Omega(F) \geq \int U_p^{m^*}(\theta) dF(\theta),$$

for all $F \in \mathcal{F}$, condition (5) is trivially sufficient for F^* to minimize Ω on $\mathcal{F}$. ■

In the next subsections, we analyze the best responses of this zero-sum game in order to find a robust mechanism.

3.1 Principal's best response

For $m^* \in \mathcal{M}$, let

$$\theta^* \in \arg\min_{\theta\in[0,1]} U_p^{m^*}(\theta)$$

be the worst state realization for the principal. By Lemma 5 in Appendix A, such state exists.[7] Moreover, by the definition of θ^* we clearly have that

$$\int U_p^{m^*}(\theta)dF^*(\theta) \geq U_p^{m^*}(\theta^*). \tag{6}$$

Suppose that condition (5) holds for mechanism $m^* \in \mathcal{M}$ and distribution $F^* \in \mathcal{F}$. Taking $dF = \delta_{\theta^*}$ the Dirac distribution which places all weight on θ^*, and combining conditions (5) and (6), we have

$$\int U_p^{m^*}(\theta)dF(\theta) = U_p^{m^*}(\theta^*) \geq \int U_p^{m^*}(\theta)dF^*(\theta) \geq \int U_p^{m^*}(\theta)dF(\theta),$$

for all $F \in \mathcal{F}$. This last condition immediately implies that $U_p^{m^*}(\theta)$ must be constant over the support of F^*. In the equilibrium of the zero-sum game, the principal will choose a mechanism to make nature indifferent among all states with positive likelihood. From Lemma 2, the constancy of the principal's ex-post payoff is a necessary condition for a robust decision-making mechanism.

Using Lemma 1 (i), we can substitute the integral form of the agent's utility into (3) to obtain

$$U_p^{m^*}(\theta) = U_a^{m^*}(0) + 2\int_0^\theta (\overline{a}^*(\widetilde{\theta}) - y_a(\widetilde{\theta}))y_a'(\widetilde{\theta})d\widetilde{\theta} - 2b(\theta)(\overline{a}^*(\theta) - y_a(\theta)) - b(\theta)^2, \tag{7}$$

where $m^* = (\overline{a}^*, \sigma^{2*})$.

If $m^* = (\overline{a}^*, \sigma^{2*}) \in \mathcal{M}$ is incentive compatible, Lemma 1 (ii) ensures that the optimal expected decision $\overline{a}^*(\theta)$ is non-decreasing and consequently differentiable almost everywhere. Imposing that $U_p^{m^*}(\theta)$ is constant over all states in $[0,1]$, differentiating condition (7) with respect to θ and equating to zero, we obtain the following ODE:

$$\frac{d\overline{a}^*}{d\theta}(\theta) = \frac{y_p'(\theta)}{b(\theta)}(\overline{a}^*(\theta) - y_p(\theta)), \tag{8}$$

for almost all $\theta \in [0,1]$.

Up to the initial condition $\overline{a}^*(0)$, the ODE (8) completely pins down the expected decision to be made by the principal in a robust mechanism.[8]

It is clear that, for any $F \in \mathcal{F}$ nature might choose, the best the principal can do is to design a mechanism satisfying the ODE (8) for some initial condition $\overline{a}^*(0)$. Notice that the solution of this ODE with initial condition $y_p(0) \leq \bar{a}^*(0)$ is non-decreasing and has increasing rate lower than or equal to one of the principal's favorite decision. Since the bias function is non-decreasing, if $\bar{a}^*(0) \leq y_a(0)$, then $\bar{a}^*(\theta) \leq y_a(\theta)$, for all $\theta \in [0,1]$. By Lemma 1, choosing $\sigma^{*2}(0) = 0$, the mechanism m^* is incentive-compatible. This discussion proves:

Proposition 1 *Let $(\overline{a}^*, \sigma^{2*}) \in \mathcal{M}$ be the mechanism of a Nash equilibrium of the zero-sum game such that the distribution has full support on $[0,1]$. Then, $\overline{a}^*(\theta)$ must satisfy the ODE (8). Moreover, any expected decision $\bar{a}^*(\theta)$ satisfying the ODE (8) with initial condition $\bar{a}^*(0) \in [y_p(0), y_a(0)]$ is incentive-compatible.*

[7]Indeed, as $U_p^{m^*}(\theta)$ is lower semi-continuous and $[0,1]$ is compact, by Weierstrass Theorem, there is a solution for $\min_{\theta\in[0,1]} U_p^{m^*}(\theta)$.

[8]It is worth noting that the only restriction that the indifference condition imposes on the distribution $P_\theta^*(\cdot)$ is on its first moment, $\overline{a}^*(\theta)$. Since $\overline{a}^*(\theta)$ is differentiable a.e. and coincides with the unique continuously differentiable solution of the ODE (8), $\overline{a}^*(\theta)$, it must be continuously differentiable as well.

The expected decision characterized by the ODE (8) ensures that principal's payoff is constant over all state realization. Hence, the solution of (8) makes the nature indifferent among all distributions in $\mathcal{F}$ and has fairly natural interpretation. Dividing both sides of (8) by $y_p'(\theta)$, this condition reads

$$\frac{d\bar{a}^*}{dy_p} = \frac{\bar{a}^* - y_p}{b},$$

which means that the derivative of the expected decision with respect to the principal's favorite decision is the ratio of two bias: the expected decision and the agent's favorite decision to the principal's favorite decision. By the discussion just before Proposition 1, if $\bar{a}^*(0) \in [y_p(0), y_a(0)]$, then this derivative is lower than or equal to one, meaning that the expected decision responsiveness to the principal's favorite decision is weakly lower than one.

This first indifference principle built in the ODE (8) characterizes the principal's best-response.

3.2 Nature's best response

In subsection 3.1, we showed that the principal chooses the decision-making mechanism to leave nature indifferent among all states it might choose. We now derive what is the distribution F^* to be chosen by nature so to induce the principal to behave as in Proposition 1.

Toward that, we start by assuming that F^* is continuously differentiable on $(0,1)$ with possibly a mass point only at $\theta = 0$ – we, of course, later verify that this is indeed the case. Using Lemma 1 (i) and integration by parts, the principal's ex-ante payoff under the distribution F^* and a robust mechanism $m^* = (\overline{a}^*, \sigma^{2*}) \in \mathcal{M}$ can be written as[9]

$$\begin{aligned} \int_0^1 U_p^m(\theta)dF^*(\theta) &= (1 - F^*(0))U_a^m(0) + F^*(0)U_p^m(0) \\ &+ \int_0^1 \{2(\overline{a}^*(\theta) - y_a(\theta))[y_a'(\theta)(1 - F^*(\theta))d\theta - b(\theta)dF^*(\theta)] - b(\theta)^2 dF^*(\theta)\}. \end{aligned} \tag{9}$$

Given the linearity of the principal's ex-ante payoff in $\overline{a}$ the only way to reconcile the optimality of the mechanism with expected decision characterized by Proposition 1 is if the nature makes the principal indifferent among all incentive-compatible mechanisms which give the same payoff at state $\theta = 0$. Therefore, the nullity of the term within the brackets of the integrand in the expression (9), i.e.,

$$\frac{dF^*(\theta)}{1 - F^*(\theta)} = \frac{y_a'(\theta)}{b(\theta)}, \tag{10}$$

joint with condition $F^*_-(1) = 1$ make the principal indifferent among all possible mechanisms in $\mathcal{M}$ with the same initial condition $\bar{a}^*(0)$.

The solution of the ODE (10) leads to the distribution of a Nash equilibrium of our zero-sum game.

Proposition 2 *Let $F^* \in \mathcal{F}$ be the distribution of a Nash equilibrium of the zero-sum game with support on $[0,1]$ such that $F^*_-(1) = 1$. Then F^* must solve the ODE (10) and, consequently,*

$$F^*(\theta) = 1 - \exp\left(-\int_0^\theta \frac{y_a'(\widetilde{\theta})}{b(\widetilde{\theta})}d\widetilde{\theta}\right)\left[1 - \exp\left(-\int_\theta^1 \frac{y_a'(\widetilde{\theta})}{b(\widetilde{\theta})}d\widetilde{\theta}\right)\right].$$

[9] In Appendix B, we derive the expression related to the integration by parts.

The solution of (10) makes the principal indifferent among all mechanisms in $\mathcal{M}$. Dividing both sides of (10) by $y_a'(\theta)$, this condition reads

$$\frac{d\ln(1-F^*)}{dy_a} = -\frac{1}{b},$$

which means that the derivative of the log-likelihood of high states with respect to the agent's favorite decision is the reciprocal inverse of the bias. Notice that the distribution F^* may have a mass point at state $\theta = 0$. The interpretation is that, given the indifference condition (10), the nature can "punish" the principal only at the lowest state. In such state, the nature chooses the maximum weight consistent with a probability distribution satisfying the indifference condition (10). This lowest state is the least favorable state for the principal once the bias is the lowest among all possible states.

This second indifference principle given by the ODE (10) characterizes the nature's best response.

3.3 Robust mechanisms

By Lemma 2, the mechanism characterized in Proposition 1 and the distribution in Proposition 2 are, respectively, the robust decision making mechanism and the worst-case distribution of states. Putting these results together we get:

Proposition 3 *A Nash equilibrium of the zero-sum game,* $(\overline{a}^*, \sigma^{2*}, F^*) \in \mathcal{M} \times \mathcal{F}$*, is the distribution* F^* *given by Proposition 2 and a mechanism with: (i) expected decision* $\bar{a}^*$ *given by Proposition 1 with* $\bar{a}^*(0) = y_a(0) - F^*(0)b(0)$ *; and (ii) variance* σ^{2*} *given by Lemma 1 (iii) with* $\sigma^{2*}(0) = 0$*.*

It is interesting to notice that the robust mechanism in Proposition 1 stems from the principal's desire to insure against her uncertainty regarding the distribution of states. Curiously, such insurance takes a very stark form: the principal designs a mechanism that assures her the same payoff for all states that might occur with positive likelihood. Put differently, an important property of a robust decision-making mechanism is the guarantee that the principal will do equally well in terms of her payoff whatever state realizes.

We now move to derive additional properties of a robust decision-making process. We split the analysis in two, depending on wether the bias function, $b(\theta)$, is constant or not. Before proceeding, however, we pause to talk about indirect implementation.

Since the influential work of Holmström (1977, 1984), a significant body of the literature on decision-making debates whether decision-making mechanisms can be implemented by delegation sets. A delegation set mechanism is an indirect mechanism in which the principal designs a set of decisions $\mathcal{D}$ and then allows the agent to make any decision as long as it lies in $\mathcal{D} \subseteq \mathcal{A}$.

For our model, the next lemma establishes a simple necessary and sufficient condition for implementability through delegation sets.

Lemma 4 *[Implementation through delegation sets] The robust (direct) decision-making mechanism* $(\overline{a}^*, \sigma^{2*})$ *can be implemented through a delegation set if and only if the mechanism is deterministic, i.e.,* $\sigma^{2*}(\theta) = 0$*, for all* $\theta \in [0, 1]$*.*

The sufficiency part of Lemma 4 can be interpreted as the counterpart of Taxation Principle (see Salanié, 1997) for contracting problems in which side payments among the agent and the principal are not allowed. Holmström (1977, 1984) and Alonso and Matouschek (2008), who, from the outset, restrict attention to deterministic mechanisms, have established a similar sufficiency result for their decision-making problems with single privately informed agents. The necessity part is straightforward: a stochastic mechanism calls for a decision-process that induces a non-degenerate distribution over $\mathcal{A}$. A mechanism with a delegation set, however, forces the agent to pick a single decision.

In the process of deriving more properties of the robust decision-making mechanism we will also check if the robust decision-making process can be implemented through delegations sets for the two cases of interest.

Constant bias

Let us exceptionally suppose that the set of states and agent's favorite decisions are unbounded, i.e., $\theta \in [0, \infty)$ and $\lim_{\theta \to \infty} y_a(\theta) = \infty$ (in particular, $\mathcal{A} = [\underline{a}, \infty)$).[10] In our model, the case in which the bias function is constant, i.e., $b(\theta) = b$, for all $\theta \in [0, \infty)$, corresponds to the main example in the seminal paper of Crawford and Sobel (1981). The differences from that paper are two. The first is that the principal can commit to a mechanism in our model. The second is that the principal is uncertain about the distribution of states. Those two differences combined lead to the following result:

Proposition 4 *[Full delegation] If set of states and agent's favorite decisions are unbounded, and the bias function is constant, then the robust decision-making mechanism has* $\overline{a}^*(\theta) = y_a(\theta)$ *and* $\sigma^{2*}(\theta) = 0$*, for all* $\theta \in [0, \infty)$*. This mechanism can be indirectly implemented through a full delegation mechanism, that is, through a delegation set equals* $\mathcal{A}$*.*

Proposition 4 is a joint implication of principal's ability to commit and her uncertainty about the distribution of states. Indeed, it is straightforward that a DM without commitment would never pick – or induce, in case one thinks of the (indirect) implementation by the delegation set $\mathcal{A}$ – the agent's most preferred decision in all states. As Crawford and Sobel (1981) first showed, a principal without commitment would choose the expected value of her most preferred decision given her information set. Such policy does not coincide with the one that agent would adopt.

It has also been extensively shown (a few examples are Melumad and Shibano (1997), Kováč and Mylanov (2008) and Alonso and Matouschek (2008)) that a principal who can commit any mechanism and knows the distribution of states would never find it optimal to fully delegate the decision to the agent, if his bias is sufficiently large.[11] In fact, for an agent with large positive bias $b > 0$, capping the delegation set at a decision $\overline{d} < y_a(1)$ is usually optimal. In contrast, for the case in which the distribution is unknown, an attempt to cap the decision at $\overline{d}$ will be immaterial as the nature has freedom to move weights around from different states. Hence, when confronted with an adversary nature and facing an agent with a constant bias, the only way the principal can assure a constant payoff across unbounded states (which is her optimal way to insure against uncertainty) is by giving full discretion to the agent.

[10] Notice that the definition of direct mechanism naturally extends to this case as $m := (\overline{a}, \sigma^2) : [0, \infty) \to \mathcal{A} \times \mathbb{R}_+$ integrable function. Also, all the results previously derived can be easily adapted to this case as well.

[11] When the bias is small, full delegation might be optimal in the certainty case when the principal can commit to mechanisms. As an example, for sufficiently small bias (or when, using their terminology, the agent is moderate), full delegation is optimal for Alonso and Matouschek's (2008) normal-linear case.

Non-constant bias

For the case in which the bias function is non-constant, matters are slightly more complicated. Our next result shows that the robust mechanism is stochastic.

Proposition 5 *If $b(\theta)$ is a non-constant function, the robust decision-making mechanism is stochastic (i.e., $\sigma^{2*}(\theta) > 0$, for some θ). Hence, it can never be implemented through a delegation set.*

Under mild regularity conditions on the distribution of states and the bias function $b(\theta)$, Kováč and Mylanov (2008) show that, for settings such as ours in which players have quadratic preferences, the optimal mechanism is deterministic when the principal knows the distribution. In contrast to the certainty case, Proposition 5 shows that randomness in decision-making processes plays a key role under uncertainty.

In such case, one may interpret the decision process as being contingent on the realization of a (payoff irrelevant) random device. Such device, somewhat surprisingly, plays the role of guaranteeing that the principal is fully insured against uncertainty when the degree of misalignment with the agent varies with the state. Put differently, the best way the principal finds to insure against the opportunistic behavior of an adversary nature is by forcing nature to face itself some "uncertainty" implied by a random decision-making process.

Although a full characterization of the robust decision-making mechanism depends on the specific form of the bias $b(\theta)$ – which, in turn, depends on $y_a(\theta)$ and $y_p(\theta)$ –, we are able to establish a few general properties of the expected decision, $\overline{a}^*(\theta)$, for the case in which $b'(\theta) > 0$ for all θ. In particular, there is no set of states for which the bias is constant.

Proposition 6 *If $b'(\theta) > 0$, for all θ, then $\overline{a}^*(\theta) \in (y_p(\theta), y_a(\theta)]$ and $\frac{d\overline{a}^*}{d\theta}(\theta) > 0$, for almost all θ.*

The first part of Proposition 6 states that the expected decision is strictly between the principal's and the agent's favorite ones with probability one. As for the second part – which says that the expected decision is always responsive to the state –, one implication is that, in contrast to the case in which the distribution is known, it is never optimal to "cap" the agent. As intuition might suggest, the degree of responsiveness of $\overline{a}^*(\theta)$ to the state θ will be larger the greater the responsiveness of the principal's favorite decision to the state, $y_p'(\theta)$, on the one hand, and smaller the bias $b(\theta)$ on the other hand.

4 Concluding remarks

We have considered the problem faced by a DM who has to rely on an informed (but biased) agent to make a decision, is uncertain about the distribution from which the relevant state is drawn and, in face of this uncertainty, has a maximin objective. The results were summarized and discussed in the introduction. We, therefore, conclude with some avenues for future research.

As opposed to what we have assumed, it is often the case (in organizations, for instance) that the DM can count on multiple agents as sources of information about the state. Thus, a natural extension of our model is, in the spirit of Martimort and Semenov (2008), to consider robust decision-making with multiple agents. Allowing for side payments between the principal and the agent, as in Krishna and Morgan (2008), is yet another interesting possibility, much as considering the possibility of repeated interactions between the principal and the agent (and a "robust learning" procedure by the former) as Alonso and Matouschek (2007). Hopefully future research will address these questions.

Appendix A: omitted proofs

Proof. (Lemma 1) The envelope theorem and usual arguments show that conditions (i) and (ii) are necessary for incentive compatibility. The condition (iii) is the necessary non-negativity of the variance. For the sufficiency, take a mechanism $(dP_\theta(\cdot))_{\theta\in[0,1]} \subset \mathbb{P}(\mathcal{A})$ that has mean and variance satisfying conditions (i)-(iii). Hence, incentive compatibility is equivalent to

$$\begin{aligned} U_a^{(\overline{a},\sigma^2)}(\theta) &\geq -\int_{\mathcal{A}} (a - y_a(\theta))^2 \, dP_{\hat{\theta}}(a) \\ &= -(\overline{a}(\hat{\theta}) - y_a(\theta))^2 - \sigma^2(\hat{\theta}) \\ &= U_a^{(\overline{a},\sigma^2)}(\hat{\theta}) - 2(y_a(\hat{\theta}) - y_a(\theta))(\overline{a}(\hat{\theta}) - y_a(\hat{\theta})) - (y_a(\hat{\theta}) - y_a(\theta))^2, \end{aligned}$$

where we used (ED), (V) and the definition of the rent function $U_a^{(\overline{a},\sigma^2)}(\theta)$ at the second and third lines. Using condition (i), these last inequalities are equivalent to

$$2(y_a(\hat{\theta}) - y_a(\theta))(\overline{a}(\hat{\theta}) - y_a(\hat{\theta})) + (y_a(\hat{\theta}) - y_a(\theta))^2 + 2\int_{\widehat{\theta}}^{\theta} (\overline{a}\left(\widetilde{\theta}\right) - y_a\left(\widetilde{\theta}\right)) y_a'\left(\widetilde{\theta}\right) d\widetilde{\theta} \geq 0.$$

For each $\hat{\theta}$, the left-hand side of the last inequality is zero for $\theta = \widehat{\theta}$ and its derivative with respect to θ is given by

$$-2(\overline{a}(\hat{\theta}) - y_a(\hat{\theta}))y_a'(\theta) - 2(y_a(\hat{\theta}) - y_a(\theta))y_a'(\theta) + 2(\overline{a}(\theta) - y_a(\theta))y_a'(\theta) = 2(\overline{a}(\theta) - \overline{a}(\hat{\theta}))y_a'(\theta)$$

which is positive if and only if $\overline{a}(\theta) - \overline{a}(\hat{\theta}) \geq 0$, once $y_a'(\theta) > 0$. By condition (ii), this is the case if and only if $\theta \geq \hat{\theta}$. Then, incentive compatibility is ensured.

Finally, we have to build a mechanism that has the given mean and variance. Define the following mechanism:

$$dP_\theta(\cdot) = \frac{1}{2}\left(\delta_{\overline{a}(\theta)-\sigma(\theta)}(\cdot) + \delta_{\overline{a}(\theta)+\sigma(\theta)}(\cdot)\right)$$

for each $\theta \in [0,1]$, where δ_x is the Dirac measure concentrated at x. It is easy to see that the mean and variance of such mechanism are exactly $\overline{a}(\theta)$ and $\sigma^2(\theta)$.[12]

By item (ii), $\overline{a}(\theta)$ is differentiable a.e. This condition and item (i) imply that $\frac{d}{d\theta}\sigma^2(\theta) = -2\frac{d}{d\theta}\overline{a}(\theta)(\overline{a}(\theta) - y_a(\theta))$, $\theta \in [0,1]$ a.e. Hence, $\frac{d}{d\theta}\sigma^2(\theta) \geq 0$, $\theta \in [0,1]$ a.e. since $\overline{a}(\theta) \leq y_a(\theta)$. This condition and item (iii) imply that $\sigma^2(\theta) \geq 0$ if and only if $\sigma^2(0) = -(\overline{a}(0) - y_a(0))^2 - U_a^{(\overline{a},\sigma^2)}(0) \geq 0$. ■

Lemma 5 *For any incentive compatible mechanism $m \in \mathcal{M}$, the principal's payoff, $U_p^m(\theta)$, is a lower semi-continuous function of θ.*

Proof. (Lemma 5) The principal's ex-post payoff at a mechanism $(\overline{a}, \sigma^2)$ reads

$$\begin{aligned} U_p^{(\overline{a},\sigma^2)}(\theta) &= \int_{\mathcal{A}} [u_a(a,\theta) - 2b(\theta)(a - y_a(\theta)) - b(\theta)^2] dP_\theta(a) \\ &= u_a(\overline{a}(\theta), \theta) - \sigma^2(\theta) - 2b(\theta)(\overline{a}(\theta) - y_a(\theta)) - b(\theta)^2. \end{aligned}$$

[12] Notice that there may exist many other mechanisms that have the same mean and variance. Also, we are implicitly using the positiveness of $\sigma^2(\theta)$ guaranteed by condition (iii) in the definition of such mechanism.

In any incentive compatible mechanism $(\overline{a}, \sigma^2)$,

$$U_a^{(\overline{a},\sigma^2)}(\theta) = u_a(\overline{a}(\theta), \theta) - \sigma^2(\theta) = U_a^{(\overline{a},\sigma^2)}(0) + 2\int_0^{\theta} (\overline{a}(\widetilde{\theta}) - y_a(\widetilde{\theta}))y_a'(\widetilde{\theta})d\widetilde{\theta}.$$

Plugging this last expression into $U_p^{(\overline{a},\sigma^2)}(\theta)$, we get

$$U_p^{(\overline{a},\sigma^2)}(\theta) = U_a^{(\overline{a},\sigma^2)}(0) + 2\int_0^{\theta} (\overline{a}(\widetilde{\theta}) - y_a(\widetilde{\theta}))y_a'(\widetilde{\theta})d\widetilde{\theta} - 2b(\theta)(\overline{a}(\theta) - y_a(\theta)) - b(\theta)^2.$$

By assumption, $b(\theta)^2$ is continuous and the term

$$U_a^{(\overline{a},\sigma^2)}(0) + 2\int_0^{\theta} (\overline{a}(\widetilde{\theta}) - y_a(\widetilde{\theta}))y_a'(\widetilde{\theta})d\widetilde{\theta}$$

is absolutely continuous and, therefore, continuous in θ. Hence, all that is left is to show that

$$-2b(\theta)(\overline{a}(\theta) - y_a(\theta))$$

is lower semi-continuous in θ. Now, both $b(\theta)$ and $y_a(\theta)$ are continuous. Hence, the above expression is lower semi-continuous if and only if $\overline{a}(\theta)$ is lower semi-continuous. Incentive compatibility implies that $\overline{a}(\theta)$ is non-decreasing. As such, it can have at most a countable number of jump discontinuities. Over the set of states for which $\overline{a}(\theta)$ is continuous, it is also lower semi-continuous. Over the set for which $\overline{a}(\theta)$ is not continuous, we can, without any loss of generality, take it to be left-continuous, which immediately implies lower semi-continuity. ■

Proof. (Proposition 3) This an immediate consequence of Lemma 1, Propositions 1 and 2, except the optimal choice of $\bar{a}^*(0)$. By the expression of the principal's expected payoff, the principal is indifferent to all choices of the expected decision at every state, except at state $\theta = 0$. For such state, the expected decision must maximize the payoff $(1 - F^*(0))U_a^{m^*}(0) + F^*(0)U_p^{m^*}(0)$, subject to $U_a^{m^*}(0) \leq -(\bar{a}^*(0) - y_a(0))^2$ (by Lemma 1 (iii)). We then have to solve the following problem:

$$\begin{aligned} \max_{\bar{a}^*(0), U_a^{m^*}(0)} \quad & U_a^{m^*}(0) - 2F^*(0)b(0)\left(\bar{a}^*(0) - y_a(0) + \tfrac{1}{2}b(0)\right) \\ s.t.\ & U_a^{m^*}(0) \leq -(\bar{a}^*(0) - y_a(0))^2. \end{aligned}$$

which is equivalent to

$$\max_{\bar{a}^*(0)} \; -(\bar{a}^*(0) - y_a(0))^2 - 2F^*(0)b(0)(\bar{a}^*(0) - y_a(0)).$$

The necessary and sufficient first-order condition gives

$$\bar{a}^*(0) = y_a(0) - F^*(0)b(0).$$

■

Proof. (Lemma 4) Necessity is a direct implication of the fact that a delegation mechanism calls for a single decision being made at each θ.

For sufficiency, define the following delegation set $\mathcal{D} = \overline{a}^*([0,1]))$. When confronted with the delegation set $\mathcal{D}$, the agent, upon observing state θ, will choose a decision to solve

$$\max_{d\in\mathcal{D}} -\left(d - y_a(\theta)\right)^2.$$

Since the direct mechanism $\left(\overline{a}^*(\theta), \sigma^{*2}(\theta)\right)_{\theta\in[0,1]}$ is non-stochastic, so that $\sigma^{*2}(\theta) = 0$ for all $\theta \in [0,1]$, incentive compatibility is equivalent to

$$-\left(\overline{a}^*(\theta) - y_a(\theta)\right)^2 \geq -\left(\overline{a}^*(\widehat{\theta}) - y_a(\theta)\right)^2, \text{ for all } \widehat{\theta} \in [0,1].$$

The above inequalities then correspond to

$$-\left(\overline{a}^*(\theta) - y_a(\theta)\right)^2 \geq -\left(d - y_a(\theta)\right)^2, \text{ for all } d \in \mathcal{D}.$$

Noticing that $\overline{a}^*(\theta) \in D$, we have that

$$\overline{a}^*(\theta) = \max_{d\in\mathcal{D}} -\left(d - y_a(\theta)\right)^2,$$

and the result is proved. ■

Proof. (Proposition 4) In the case of constant bias, unbounded states and agent's favorite decisions, it is easy to show that the distribution of Proposition 2 is given by

$$F^*(\theta) = 1 - \exp\left(-\frac{y_a(\theta) - y_a(0)}{b}\right)$$

where $F^*(0) = 0$. From Proposition 3, it suffices to show that $\overline{a}^*(\theta) = y_a(\theta)$ satisfies the ODE (8), since the initial condition is $\bar{a}^*(0) = y_a(0) - F^*(0)b = y_a(0)$.

For the case of a constant bias, $y_a(\theta) - y_p(\theta) = b(\theta) = b$, for all $\theta \in [0,\infty)$. Hence, $y_p'(\theta) = y_a'(\theta)$, for all $\theta \in [0,\infty)$. Using these facts and plugging $\overline{a}^*(\theta) = y_a(\theta)$ in the ODE (8), we have

$$\frac{d\overline{a}^*}{d\theta}(\theta) = y_a'(\theta) = \frac{y_p'(\theta)}{b}\left(y_a(\theta) - y_p(\theta)\right) = y_p'(\theta).$$

So, the ODE (8) is satisfied with $\overline{a}^*(\theta) = y_a(\theta)$. ■

Proof. (Proposition 5) We only have to prove that $\sigma^*(\theta) > 0$, for all $\theta \in [0,1]$. Suppose that $\sigma^*(\theta) = 0$, for some $\theta \in (0,1)$. We have that

$$\theta \in \arg\min_{\hat{\theta}\in[0,1]} \sigma^{*2}(\hat{\theta}) = -(\overline{a}^*(\hat{\theta}) - y_a(\hat{\theta}))^2 - U_a^{(\bar{a}^*,\sigma^{2*})}(\hat{\theta})$$

and the first-order condition gives

$$0 = \frac{d}{d\theta}\left(\sigma^{*2}(\widehat{\theta})\right)\Bigg|_{\widehat{\theta}=\theta} = -2(\overline{a}^*(\theta) - y_a(\theta))\frac{d\overline{a}^*}{d\theta}(\theta).$$

Thus, $\frac{d\overline{a}^*}{d\theta}(\theta) = 0$ or $\overline{a}^*(\theta) = y_a(\theta)$. By (8), these are equivalent to $\overline{a}^*(\theta) = y_p(\theta)$ or $\overline{a}^*(\theta) = y_a(\theta)$. If $\overline{a}^*(\theta) = y_p(\theta)$, since $y_p'(\theta) > 0 = \frac{d\overline{a}^*}{d\theta}(\theta)$, by continuity $y_p'(\theta+\epsilon) > \frac{d\overline{a}^*}{d\theta}(\theta+\epsilon)$, for every $\epsilon > 0$ sufficiently small. This implies that $\overline{a}^*(\theta+\epsilon) < y_p(\theta+\epsilon)$ and, consequently, $\frac{d\overline{a}^*}{d\theta}(\theta+\epsilon) < 0$ from (8). This last condition contradicts incentive compatibility because of Lemma 1 (ii). If $\overline{a}^*(\theta) = y_a(\theta)$, (8) implies that

$\frac{d\overline{a}^*}{d\theta}(\theta) = y_p'(\theta)$. However, this last condition contradicts Assumption A which says that $b(\theta) = y_a(\theta) - y_p(\theta)$ is increasing.

Since the derivative of $\sigma^{*2}(\hat{\theta})$ is not zero on $(0,1)$, it is always positive or negative on $(0,1)$ (see the Darboux Theorem in Protter, 1998). By the above expression of its derivative, this condition depends on whether $\overline{a}^*(0) - y_a(0) < 0$ or $\overline{a}^*(0) - y_a(0) > 0$. ■

Proof. (Proposition 6) It is easier to start the proof by the second part of the result.

Lemma 1 (ii) implies that $\frac{d\overline{a}^*}{d\theta}(\theta) \geq 0$. If $\frac{d\overline{a}^*}{d\theta}(\theta) = 0$, then $\overline{a}^*(\theta) = y_p(\theta)$. Since both $\frac{d\overline{a}^*}{d\theta}(\theta)$ and $y_p'(\theta) > 0$ are continuous functions, there must exist $\epsilon > 0$ such that $\overline{a}^*(\widehat{\theta}) - y_p(\widehat{\theta}) < 0$ and $\frac{d\overline{a}^*}{d\theta}(\widehat{\theta}) < 0$, for all $\widehat{\theta} \in [\theta, \theta + \epsilon)$, which violates incentive compatibility.

Incidentally, the above reasoning establishes that $\overline{a}^*(\theta) > y_p(\theta)$, for all θ. In fact, since incentive compatibility calls for $\overline{a}^*(\theta)$ being non-decreasing and, by Proposition 1,

$$\frac{d\overline{a}^*}{d\theta}(\theta) = \frac{y_p'(\theta)}{b(\theta)}(\overline{a}^*(\theta) - y_p(\theta)),$$

one must have $\overline{a}^*(\theta) \geq y_p(\theta)$. Our proof of the responsiveness result relies on the fact that one cannot have $\overline{a}^*(\theta) = y_p(\theta)$. Hence, $\overline{a}^*(\theta) > y_p(\theta)$, for all θ.

Now, we argue that there can be no open interval $\Theta = (\underline{\theta}, \overline{\theta})$ for which $\overline{a}^*(\theta) = y_a(\theta)$ if $\theta \in \Theta$. Indeed, if there were such set, one would have

$$\frac{d\overline{a}^*}{d\theta}(\theta) = y_p'(\theta) \text{ for all } \theta \in \Theta,$$

so that $\overline{a}^*(\theta) = y_a(\theta) = \overline{a}^*(\underline{\theta}) + y_p(\theta) - y_p(\underline{\theta}) = y_a(\underline{\theta}) + y_p(\theta) - y_p(\underline{\theta})$ or $b(\theta)$ is constant over Θ. ■

Appendix B: integration by parts

In this appendix, we perform the integration by parts that leads to the formula for the principal's objective in subsection 3.2. Let $m = (\bar{a}, \sigma^2)$ and suppose that f is the absolutely continuous part of the measure dF with $F_-(1) = 1$ (i.e., there is no mass point at $\theta = 1$). We have that

$$\begin{aligned}
\int_0^1 U_p^m(\theta)dF^*(\theta) &= U_a^m(0) + 2\int_0^1 \left[\int_0^\theta (\overline{a}(\widetilde{\theta}) - y_a(\widetilde{\theta}))y_a'(\widetilde{\theta})d\widetilde{\theta} - b(\theta)(\overline{a}(\theta) - y_a(\theta)) - \tfrac{1}{2}b(\theta)^2\right] dF(\theta) \\
&= U_a^m(0) + F(0)(U_p^m(0) - U_a^m(0)) + 2(F(\theta) - 1)\int_0^\theta (\overline{a}(\widetilde{\theta}) - y_a(\widetilde{\theta}))y_a'(\widetilde{\theta})d\widetilde{\theta}\Bigg|_0^1 \\
&\quad + 2\int_0^1 \left[(\overline{a}(\theta) - y_a(\theta))y_a'(\theta)\tfrac{1-F(\theta)}{f(\theta)} - b(\theta)(\overline{a}(\theta) - y_a(\theta)) - \tfrac{1}{2}b(\theta)^2\right] f(\theta)d\theta \\
&= (1 - F(0))U_a^m(0) + F(0)U_p^m(0) \\
&\quad + 2\int_0^1 \left[(\overline{a}(\theta) - y_a(\theta))\left(y_a'(\theta)\tfrac{1-F(\theta)}{f(\theta)} - b(\theta)\right) - \tfrac{1}{2}b(\theta)^2\right] f(\theta)d\theta.
\end{aligned}$$

References

[1] Alonso, R., Dessein, W. and Matouschek, N., 2007, "Relational Delegation", *RAND Journal of Economics*, vol. 38(4), pp. 1070-89.

[2] Alonso, R., Dessein, W. and Matouschek, N., 2008, "Optimal Delegation", *Review of Economic Studies*, vol. 75(1), pp. 259-93.

[3] Crawford, V. and Sobel, J., 1982, "Strategic Information Transmission", *Econometrica*, vol. 50(6), pp. 1431-51.

[4] Dessein, W., 2002, "Authority and Communication in Organizations", *Review of Economic Studies*, vol. 69(4), pp. 811-38.

[5] Frankel, A., 2013, "Aligned Delegation", *American Economic Review*, forthcoming.

[6] Homström, B., 1977, "On Incentives and Control in Organizations", Ph.D. Dissertation, Stanford University.

[7] Homström, B., 1984, "On the Theory of Delegation", in *Bayesian Models in Economic Theory,* edited by M. Boyer and R.E. Kihlstrom, vol. 5 of *Studies in Bayesian Economics*, pp. 115-41.

[8] Kováč, E., and Mylanov, T., 2009, "Stochastic Mechanisms in Settings without Monetary Transfers: The Regular Case", *Journal of Economic Theory*, vol. 144(4), pp. 1373-95.

[9] Krishna, V. and Morgan, J., 2008, "Contracting for Information under Imperfect Commitment", *RAND Journal of Economics*, vol. 39(4), pp. 905-25.

[10] Martimort, D. and Semenov, A., 2008, "The Informational Effects of Competition and Collusion in Legislative Politics", *Journal of Public Economics*, vol. 92(7), pp. 1541-63.

[11] Melumad, N., and Shibano, T., 1991, "Communication in Settings with No Transfers", *RAND Journal of Economics,* vol. 22(2), pp. 173-98.

[12] Milgrom, P. and Segal, I., 2002, "Envelope Theorems for Arbitrary Choice Sets", *Econometrica,* vol. 70(2), pp. 583-601.

[13] von Neumann, J., and Morgenstern, O., 1944, *Theory of Games and Economic Behavior.* Princeton University Press, Princeton.

[14] Osborne, M., and Rubinstein, A., 1994, *A Course in Game Theory.* The MIT Press, Cambridge, MA.

[15] Protter, M., 1998, *Basic Elements of Real Analysis.* Springer-Verlag, New York.

[16] Salanie, B., 1997, *The Economics of Contracts.* The MIT Press, Cambridge, MA.