TEXTO PARA DISCUSSÃO

No. 642

Robust Mechanisms: the curvature case

Vinicius Carrasco
Vitor Farinha Luz,
Paulo Monteiro
Humberto Moreira

DEPARTAMENTO DE ECONOMIA
www.econ.puc-rio.br

Robust Mechanisms: the curvature case[*†]

Vinicius Carrasco, Vitor Farinha Luz, Paulo Monteiro and Humberto Moreira[‡]

First Version: June 2015[§]

Abstract

This note considers the problem of a principal (she) who faces a privately informed agent (he) and only knows one moment of the distribution from which his types are drawn. Payoffs are non-linear in the allocation and the principal maximizes her worst-case expected profits. We recast the robust design problem as a zero-sum game played by the principal and an adversarial nature who seeks to minimize her expected payoffs. The robust mechanism and the worst case distribution are, then, the Nash equilibrium of such game. A robustness property of the optimal mechanism imposes restrictions on the principal's ex-post profit function. These restrictions then lead to the optimal mechanism. The robust mechanism entails exclusion of low types and distortions at the intensive margin that (in a precise sense) are larger than what those that prevail in standard Bayesian mechanism design problems.

Keywords: Robust Mechanism Design, Monopolistic Screening under Uncertainty, Taxation and Regulation under Uncertainty, Dynamic and Multidimensional Robust Design

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

---

[*]This note is being extended to incorporate two sets of analysis: (i) a dynamic version of the main model in which the agent's information evolves over time and all is known by the principal is that types satisfy a martingale condition, and (ii) a multidimensional version of the model. We are also working on the full derivation of examples (e.g., regulation and taxation) that use the results this work derives. A paper incorporating (i), (ii) and the examples will subsume this note.

[†]We have benefited from conversations with Gabriel Carroll, Nicolás Figueroa, Stephen Morris, Leonardo Rezende and Yuliy Sannikov. Moreira gratefully acknowledges financial support from CNPq.

[‡]Carrasco: Department of Economics, PUC-Rio, vncarrasco@gmail.com. Luz: European University Institute and the Department of Economics, The University of British Columbia, vitor.farinha@gmail.com. Monteiro: FGV/EPGE, PKLM@fgv.br. Moreira: FGV/EPGE, humberto@fgv.br.

[§]This work encompasses a previous paper, by a subset of the current authors, that has been presented in several occasions since December 2012.

1 Introduction

In search for predictions for contracting models that do not rely on fine details of their primitives – which, in practice, are unlikely to be known by real world designers – a growing literature assumes that a principal (she) might not be fully aware of the distribution form which an agent's (he) private information is drawn, e.g., Bergemann and Schlag (2008), Carroll (2013), Carrasco et al. (2015).[1] One common feature of all such papers is that they consider the case in which the agent's and the principal's ex-post payoffs are linear in the allocation. This is a drawback, as, in a wide variety of applications – such as optimal taxation, screening, decision-making and optimal regulation, among others –, its is natural to assume some curvature in payoffs: e.g., for some insurance to be desirable in a taxation setting, marginal utilities must decrease with consumption; the further away a decision is from an agent's favorite one, the worse off such agent is; effort towards cost reduction in a regulatory setting becomes costlier the more the agent exerts it. In this paper, we fill this gap by considering a single agent mechanism design problem in which, first, payoffs are non-linear in the allocation, and, second, the principal knows only the first moment of the distribution from which the agent's types is drawn. In face of such uncertainty, maximizes her worst case expected profits. We call the mechanism that results of such procedure the "robust mechanism".

The environment we consider is general enough to have, as particular cases, Mussa and Rosen (1978)'s screening setting, Laffont and Tirole (1986)'s regulation model and Mirrlees (1971)'s taxation model. To assume that the principal has some knowledge of the distribution is key to make the problem interesting for the cases in which leaving information rents to the agent is costly. In fact, for such cases, if we were to assume full ignorance of the distribution from the principal's point of view, the robust mechanism would be trivial: probability one is assigned to the least favorable type from the principal's perspective, and the optimal mechanism is the complete information one.

Finding the robust mechanism for the "curvature" case is substantially harder than for the case of linear payoffs. However, one can still recast, as Bergemann and Schlag (2008) and Carrasco et al. (2015) do, the robust design problem as a zero-sum game played by the designer and an adversarial nature, who chooses distribution for the agent's private information to minimize the designer expected payoffs. A Nash equilibrium of such game is, then, the robust mechanism and the worst case distribution.

In the zero-sum game, taking as given the mechanism chosen, nature minimizes the designer expected payoff by choosing distributions that satisfy two constraints: first of all, the distribution

[1] Hurwicz and Shapiro (1978), Garrett (2014) and Frankel (2014) consider the case in which the principal is (partially) unaware of the agent's payoff.

must integrate to one; second, its expected value has to equal to some constant k. Those constraints express what the designer knows about the distribution. Once nature incorporates those constraints in the Lagrangian functional, its problem becomes one of minimizing, by choice of distributions, the designer's expected payoff subtracted by inner product of the constraints and their shadow costs (the Lagrangian multipliers). We are then able to show that will only place positive likelihood on types θ in set support of the type set $[0, 1]$ for which the designer's payoff equal $\xi\theta - \lambda$, where ξ is the shadow cost of the constraint that imposes that the average of the distribution must be k, and λ is the shadow cost of the constraint that the distribution must integrate to one. It then follows that, at a robust mechanism, the designer's payoff is piecewise linear in the agent's (reported) types, so that robustness imposes restrictions on the designer's payoff level. Using the game theoretical interpretation of robustness, one has that, by picking a mechanism that induces piecewise linear payoffs, the designer assures that nature will be indifferent among all feasible distributions. Since it is indifferent, nature is willing to choose a distribution that guarantees the designer will find it optimal to design a mechanism that solves the ordinary differential equation (ODE). We derive such distribution and, therefore, fully characterize the Nash equilibrium of the zero-sum game.

In general, one cannot derive a closed form solution for the ODE that defines the robust scheme, but general properties of optimal allocation can be derived, and their implications for the examples we consider discussed. First of all, in all design problems in which the principal finds it costly to leave informational rents for the agent, low types are bunched in a single allocation. In particular, for the Laffont and Tirole (1986) regulation model, this implies that there will be an interval of types (the least efficient ones) which will face cost-plus contracts. In the Mussa and Rosen (1978) screening setting, in turn, this amounts to exclusion of consumer's will low willingness to pay. Second, in a sense that we make precise in text, robust contracts are less powerful than their Bayesian counterparts: in Laffont and Tirole (1986) context, this means that regulation contracts will always be bounded away from fixed-price arrangements; whereas, in Mussa and Rosen (1978), quantity discrimination will always take place.

Related literature

Our paper is part of a growing literature on mechanism design with principals with maxmin preferences.[2] Our work differs from Frankel (2014)'s and Carrasco and Moreira (2013)'s decision-making problems because the utility is fully transferable in our setting. This, in turn, allows for more general mechanisms than the "delegation" ones those papers derive. Garrett (2014)

[2]There is also a growing literature with maxmin agents. Bose et al. (2006) and Wolitsky (2014) are examples of analysis of, respectively, optimal auctions and bilateral trade when agents have maxmin preferences.

considers the case of a principal who does not know the producer's disutility of effort, and shows that a simple fixed-price-cost-reimbursement (FPCR) menu minimizes the principal's maximum expected payment to the agent. In Carroll (2015a), the principal only partially knows the set of actions available to the agent; he shows that if the principal maximizes expected profits under worst-case set of actions, the optimal contract is linear in output. We, in turn, assume that payoffs are common knowledge, and consider the implications of uncertainty regarding the distribution of types.

The papers that are closest to ours are the ones in robust pricing literature pioneered by Bergemann and Schlag (2008, 2011). In their first paper, they consider the case in which the seller designs a mechanism to minimize the maximum regret, whereas in the second they also consider a max min procedure. In independent work, Carroll (2013) and Carrasco et al. (2015) consider a setting in which only one moment of the distribution of willingness to pay of the consumer is known. Our paper differs from all this by considering the case of curvature in ex-post payoffs.

Much as Carrasco et al. (2015) do, we show that a variant of the model in which the agent's type evolves over time and satisfies a martingale property leads to a time-invariant robust mechanism: the period by period repetition of the static mechanism is optimal. This – along with the fact that we consider maxmin design (with the restriction that expected values follows martingale[3]) and allow for general mechanisms – is the main difference from what is obtained by Handel and Misra (forthcoming). They show, in a setting in which monopolist launches a new product and, without knowledge of (the time invariant) demand, decides – restricting attention to price posting mechanisms – on intertemporal prices to minimize maximum intertemporal regret, that prices decrease over time if consumers are homogenous, and increase if consumers are heterogeneous. Caldentey et al. (2015) also consider minimax intertemporal pricing for the case in which the seller restricts attention to posting price mechanisms. However, on top of not knowing demand, the seller does not know the arrival process of consumers in their paper. They also establish that optimal price paths are decreasing when buyers are rational, which contrasts to our time invariance result.

In a general additive model, Carroll (2015b) proves that the optimal multidimensional mechanism when an ambiguity averse designer knows the marginal distributions in each dimension, but is uncertain about the joint distribution of types, entails full separation. Carrasco et al. (2015) prove a separability result for the case in which a seller only knows the average of a consumer's multidimensional willingness to pay, but the agent's payoff is linear. In this paper, the separability result is a corollary of Carroll's result.

[3]Again, this can be justified by an information acquisition story, since, if information is acquired over time, at any given time the martingale property must hold.

Organization

In addition to this introduction, this note is composed of the next section, which lays down the model and fully derives the robust mechanism, and Section 3 with conclusions.

2 Model

There is a principal and and agent, who is is privately informed about a single dimensional parameter θ in the interval $[0, 1]$, that affects his payoff and the principal's. The principal's payoff is

$$t - c(q)$$

whereas the agent's is

$$u(q, \theta) - t$$

with $u(x, \theta)$ satisfying a (weak) single crossing condition and t is the transfer between the agent and principal. That is, we make the the following assumptions on the surplus function $s(q, \theta) := u(q, \theta) - c(q)$:

$$s_q > 0 > s_{qq} \text{ and } s_{q\theta} > 0 \tag{1}$$

$s(0, \theta) = s(q, 0) = u(0, 0)$, for all θ, q. Finally, let us assume that the function

$$q \to \frac{s_{q\theta}(q, \theta)}{s_q(q, \theta)} \text{ is strictly increasing.} \tag{2}$$

We denote by $q^{FB}(\theta) = \underset{x \geq 0}{\arg\max}\, s(x, \theta)$ the first-best allocation. The necessary and sufficient first-order condition of this first-best solution is given by

$$s_q(q^{FB}(\theta), \theta) = 0$$

for all $\theta \in [0, 1]$. By our hypothesis, $q^{FB}(\cdot)$ is strictly increasing.[4]

The principal only knows the first moment of the distribution $k \in [0, 1]$ from which types are drawn.

From the revelation principle we restrict attention on the set of direct mechanisms, defined

[4] Indeed, notice that by our assumptions and applying the implicit derivative we get

$$\frac{dq^{FB}}{dq}(\theta) = -\frac{s_{q\theta}(q^{FB}(\theta), \theta)}{s_{qq}(q^{FB}(\theta), \theta)} > 0.$$

as

$$\mathcal{M} \equiv \{m = (q,t) : [0,1] \to [0,1] \times \mathbb{R} \text{ is } \theta\text{-measurable}\}.$$

The principal trades q for transfer t and only has information about the average of the type distribution. Formally, the set of possible distributions is

$$\mathcal{F} \equiv \left\{F \in \Delta([0,1]) ; \int \theta dF(\theta) = k\right\},$$

for some $k \in [0,1]$, where $\Delta([0,1])$ is the space of all distribution on $[0,1]$.

In face of such uncertainty, she maximizes her worst case expected payoffs. Noticing that she can rely on direct mechanisms in which the agent is truthful, her problem reads:

$$\max_{m=(q,t)\in\mathcal{M}} \min_{F\in\mathcal{F}} \int [t(\theta) - c(q(\theta))] dF(\theta)$$

subject to

$$u(q(\theta),\theta) - t(\theta) \geq u(q(\theta'),\theta) - t(\theta')$$
$$u(q(\theta),\theta) - t(\theta) \geq 0.$$

As usual, incentive compatibility is equivalent to[5]

$$\mathcal{U}(\theta) = u(q(\theta),\theta) - t(\theta) = \mathcal{U}(0) + \int_0^\theta u_\theta(q(\tau),\tau)d\tau$$

and $q(\theta)$ is non-decreasing.

Define the strategy space of the principal as:

$$\mathcal{Q} = \{q : [0,1] \to [0,\infty) \text{ non-decreasing}\}.$$

Substituting the agent's indirect utility in the principal's payoff and noticing that, regardless of distributions, it is always optimum to make $\mathcal{U}(0) = 0$, her problem can be rewritten as

$$\max_{q\in\mathcal{Q}} \min_{F\in\mathcal{F}} \int \left[s(q(\theta),\theta) - \int_0^\theta s_\theta(q(\tau),\tau) d\tau\right] dF(\theta).$$

[5]We use the subindex as an alternative notation for partial derivative. For instance, u_θ represents the partial derivative of u with respect to θ.

2.1 Robust mechanism as a Nash equilibrium of a zero-sum game

It is useful this maxmin problem as a zero-sum game, played by the principal and an adversarial nature, who chooses distributions to minimize the principal's expected payoff. A Nash equilibrium of such game will then deliver the robust mechanism (the equilibrium strategy of the principal) and the worst case distribution (the equilibrium strategy of nature). To find an equilibrium of this game, it is useful to denote the allocation by its inverse function, i.e., the type assignment function. Abusing notation, we denote the type assignment function by $\theta(q)$. Let us consider the following auxiliary ODE:

$$\begin{aligned} &\tfrac{d\theta}{dq}(q) = \xi^{-1} s_q(q,\theta(q)), \text{ for all } q \in [0,q_1] \\ &\theta(q_1) = 1, \end{aligned} \tag{3}$$

where $q_1 := q^{FB}(1)$. The following result characterizes the solution of this ODE.

Lemma 2.1. *For each* $\xi > 0$, there exists a unique solution $\theta^\xi(q)$ for the ODE 3 defined in the maximal interval $[0,q_1]$, which is increasing in q and ξ. Moreover, $\underset{\xi\to 0}{lim}\theta^\xi(0) = 0$ *and* $\underset{\xi\to\infty}{lim}\theta^\xi(q) = 1$, for all $q \in [0,q_1]$.

Proof. First notice that the conventional contraction method to prove the existence and uniqueness of solution for the ODE readily applies in this case. We can easily show that this method also implies the claimed monotonically results. Let $q_0 \geq 0$ be the lower bound of the maximal interval of the solution of the ODE (3). First, notice that $\frac{d\theta^\xi}{dq}(q) = 0$, then $\frac{d^2\theta^\xi}{dq^2}(q) < 0$. Indeed, taking the total derivative of (3) with respect to q we get:

$$\frac{d^2\theta^\xi}{dq^2}(q) = \xi^{-1}\left[s_{qq}(q,\theta^\xi(q)) + s_{q\theta}(q,\theta^\xi(q))\frac{d\theta^\xi}{dq}(q)\right] < 0.$$

Therefore, $\frac{d\theta^\xi}{dq}(q) > 0$ for all $q \in (q_0,q_1]$. Let (q^n) be a sequence in $(q_0,q_1]$ converging to q_0 and let $\underline{\theta} = \lim \theta^\xi(q^n)$. Define $\theta^\xi(q_0) = \underline{\theta}$. Notice that by the argument we just made, if $\frac{d}{dq}\theta^\xi(q^n)$ converges to zero, $\frac{d^2}{dq^2}\theta^\xi(q^n)$ must converge to a strictly negative number. Hence, it is not possible that $q_0 = 0 = \underline{\theta}$ and also $q_0 > 0$ with $\underline{\theta} = 0$.

Finally, when $\xi \to \infty$, $\frac{d\theta^\xi}{dq}(\cdot)$ converges pointwisely to zero. Then, we have that $\lim_{\xi\to\infty}\theta^\xi(q) = 1$, for all $q \in [0,q_1]$, once $\theta^\xi(q_1) = 1$. On the other hand, when $\xi \to 0$, $\frac{d\theta^\xi}{dq}(q)$ converges pointwisely to infinite if $\theta^\xi(q)$ is bounded away from the first-best frontier (the inverse function of $q^{FB}(\cdot)$), which is a contradiction. Hence, $\theta^\xi(\cdot)$ converges pointwisely to the first-best frontier and, in particular, $\lim_{\xi\to 0}\theta^\xi(0) = 0$. □

Consider $\tilde{\theta}(\xi) = \theta^\xi(0)$ and let $q^\xi(\theta)$ be the inverse function of $\theta^\xi(q)$. One then has

$$\frac{dq^{\xi}}{d\theta}(\theta)s_q\left(q^{\xi}(\theta),\theta\right)=\xi, \tag{4}$$

for all $\theta \in [\tilde{\theta}(\xi),1]$. If $\Pi^{\xi}(\cdot)$ is the profit function associated to $q^{\xi}(\theta)$, then the total derivative of $\Pi^{\xi}(\theta)$ is[6]

$$\frac{d\Pi^{\xi}}{d\theta}(\theta)=\frac{dq^{\xi}}{d\theta}(\theta)s_q\left(q^{\xi}(\theta),\theta\right)=\xi, \tag{5}$$

for all $\theta \in [\tilde{\theta}(\xi),1]$. Hence,

$$\Pi^{\xi}(\theta)=\begin{cases}\xi[\theta-\tilde{\theta}(\xi)] & \text{if } \theta\in[\tilde{\theta}(\xi),1]\\ 0 & \text{otherwise,}\end{cases} \tag{6}$$

where we are implicitly defining $q^{\xi}(\theta)=0$ for all $\theta\in[0,\tilde{\theta}(\xi)]$.

The principal's expected payoff at the selling mechanism $q^{\xi}(\theta)$ satisfying (4) will be

$$\int \Pi^{\xi}(\theta)\,dF(\theta)=\int_{\tilde{\theta}(\xi)}^{1}\xi[\theta-\tilde{\theta}(\xi)]dF(\theta)=\int_{0}^{1}\xi[\theta-\tilde{\theta}(\xi)]dF(\theta)+\int_{0}^{\tilde{\theta}(\xi)}\xi[\tilde{\theta}(\xi)-\theta]dF(\theta)\geq\xi\left[k-\tilde{\theta}(\xi))\right],$$

for all $F\in\mathcal{F}$. Notice that the above inequality becomes an equality if and only if the cumulative distribution function (cdf) F has support on $[\tilde{\theta}(\xi),1]$. Therefore, we have just proved the following:

Theorem 2.1. *[Nature's best-reply] Let $\xi\geq 0$ and assume that the principal chooses $q^{\xi}\in\mathcal{Q}$ that solves (4). Then, nature's best-reply in $\mathcal{F}$ is any cdf F with support on $[\tilde{\theta}(\xi),1]$. In this case, the expected profit is given by $\int\Pi^{\xi}(\theta)dF(\theta)=\xi[k-\tilde{\theta}(\xi))]$.*

The next theorem characterizes the cdf F, with support $[\tilde{\theta}(\xi),1]$, that induces the solution of the ODE 4 as the principal's best-reply to F. For this, denote

$$\gamma(\theta)=\frac{s_{q\theta}(q^{\xi}(\theta),\theta)}{s_q(q^{\xi}(\theta),\theta)}.$$

[6] To get this expression, notice that

$$\begin{aligned}\frac{d\Pi^{\xi}(\theta)}{d\theta} &= \frac{d}{d\theta}\left[s\left(q^{\xi}(\theta),\theta\right)-\int_0^{\theta}u_{\theta}\left(q^{\xi}(\tau),\tau\right)d\tau\right]\\ &= s_q\left(q^{\xi}(\theta),\theta\right)\frac{dq^{\xi}}{d\theta}(\theta)+u_{\theta}\left(q^{\xi}(\theta),\theta\right)-u_{\theta}\left(q^{\xi}(\theta),\theta\right)\\ &= s_q\left(q^{\xi}(\theta),\theta\right)\frac{dq^{\xi}}{d\theta}(\theta).\end{aligned}$$

Notice that $\gamma(\widetilde{\theta}(\xi)) > 0$ and $\gamma(1) = \infty$. Define

$$F^{\xi}(\theta) = 1 - e^{-\int_{\widetilde{\theta}(\xi)}^{\theta} \gamma(u)du}, \text{ for all } \theta \in [\widetilde{\theta}(\xi), 1) \tag{7}$$

and $F^{\xi}(1) = 1$. Let $a := e^{-\int_{\widetilde{\theta}(\xi)}^{1} \gamma(u)du} < 1$ and $f^{\xi} = \frac{d}{d\theta}F^{\xi}$.

Theorem 2.2. *[Principal's best-reply] Fix $\xi > 0$. Let $q^{\xi} \in \mathcal{Q}$ be the solution of (4) and $F^{\xi} \in \mathcal{F}$ defined by (7) - which has absolutely continuous part given by f^{ξ} and mass a at 1. Then, q^{ξ} is the principal's best-reply to F^{ξ}.*

Proof. Let $q \in \mathcal{Q}$. Performing the usual integration by parts, we get

$$\int_0^1 \left[s(q(\theta),\theta) - \int_0^{\theta} s_{\theta}(q(\tau),\tau)d\tau \right] f^{\xi}(\theta)d\theta = \int_0^1 \left[s(q(\theta),\theta) - \frac{F^{\xi}_{-}(1) - F^{\xi}(\theta)}{f^{\xi}(\theta)} s_{\theta}(q(\theta),\theta) \right] f^{\xi}(\theta)d\theta$$

and, therefore, we have

$$\begin{aligned} \int_0^1 \left[s(q(\theta),\theta) - \int_0^{\theta} u_{\theta}(q(\tau),\tau)d\tau \right] dF^{\xi}(\theta) &= \int_0^1 \left[s(q(\theta),\theta) - \tfrac{F^{\xi}_{-}(1) - F^{\xi}(\theta)}{f^{\xi}(\theta)} s_{\theta}(q(\theta),\theta) \right] f^{\xi}(\theta)d\theta \\ &+ (1 - F^{\xi}_{-}(1)) \left(s(q(1),1) - \int_0^1 s_{\theta}(q(\tau),\tau)d\tau \right) \\ &= \int_0^1 \left[s(q(\theta),\theta) - \tfrac{1 - F^{\xi}(\theta)}{f^{\xi}(\theta)} s_{\theta}(q(\theta),\theta) \right] f^{\xi}(\theta)d\theta \\ &+ (1 - F^{\xi}_{-}(1)) s(q(1),1). \end{aligned}$$

Taking the pointwise derivative for each θ and imposing it is zero at $q(\theta)$ ensure that $q(1) = q_1$ and

$$s_q(q(\theta),\theta) = \frac{1 - F^{\xi}(\theta)}{f^{\xi}(\theta)} s_{\theta q}(q(\theta),\theta).$$

Then, by the definition of F^{ξ}, we have that

$$\frac{s_{\theta q}(q(\theta),\theta)}{s_q(q(\theta),\theta)} = \frac{f^{\xi}(\theta)}{1 - F^{\xi}(\theta)} = \gamma(\theta) = \frac{s_{\theta q}(q^{\xi}(\theta),\theta)}{s_q(q^{\xi}(\theta),\theta)},$$

for all $\theta \in [\tilde{\theta}(\xi), 1)$. By our hypothesis, we must have that $\tilde{q}(\theta) = q(\theta)$, for all $\theta \in [0,1]$. □

Notice that previous theorem does not ensure that $F^{\xi} \in \mathcal{F}$, i.e., $\int \theta dF^{\xi}(\theta) = k$. The next lemma shows that there exists $\xi^* > 0$ for which the associated pair allocation-cdf, (q^*, F^*), constitutes a Nash equilibrium.

Lemma 2.2. *Consider the cdf F^{ξ} of the previous theorem for each $\xi \geq 0$. Then, there exists $\xi^* > 0$ such that $\int \theta dF^*(\theta) = k$, where $F^* = F^{\xi^*}$. In particular, taking $q^* \in \mathcal{Q}$ solution of (4)* *for ξ^*, the pair $(q^*, F^*) \in \mathcal{Q} \times \mathcal{F}$ is a Nash equilibrium of the zero-sum game between the nature and the principal.*

Proof. Notice that

$$\begin{aligned} \int \theta dF^*(\theta) &= \int \theta f^*(\theta) d\theta + a^* \\ &= \int \left[F^*_-(1) - F^*(\theta)\right] d\theta + a^* \\ &= 1 - \int F^*(\theta) d\theta \\ &= \int_{\tilde{\theta}(\xi^*)}^1 e^{-\int_{\tilde{\theta}(\xi^*)}^{\theta} \gamma^*(u) du} d\theta. \end{aligned}$$

The integrand of the last integral goes 1 when $\xi^* \to \infty$ and goes to 0 when $\xi^* \to 0$, because $q^*(\theta)$ is bounded and converges to the first-best solution when $\xi^* \to 0$. Therefore, there must exist $\xi^* > 0$ such that $\int \theta dF^*(\theta) = k > 0$. □

We now discuss a bit the content of the results we have just proved. First of all, we have that robustness imposes restrictions on how the designer's payoff can vary with the agents (reported) types. More precisely, the designer's ex-post payoff are piecewise linear, conditional on the allocation $q^{\xi}(\theta)$ being strictly positive. The interpretation is that, knowing only the first moment of the distribution of types, the designer can only explore linearly higher agent's types; else nature could move likelihood weights in a way that reduces the designer's expected payoffs. The shape of the principal's ex-post payoffs fully pins-down the robust allocation; indeed, the allocation solves an ODE derived from the linearity of the principal's payoff conditional on positive $q^{\xi}(\theta)$.

The worst-case distribution is derived from a condition that is analogous the one that prevails is standard Bayesian mechanism design. In fact, in standard mechanism design problems, the designer equates the marginal social benefits of increasing q to the marginal cost of leaving informational rents to the agent. The worst case distribution guarantees that this condition holds at the allocation that induces a piecewise linear payoff to the principal. One striking feature of a robust design is that exclusion of an interval of types. In contrast to standard Bayesian models, this holds true regardless of the behavior of marginal social surplus at the zero allocation, $\frac{\partial \theta^{\xi}}{\partial q}(0)$.[7] By excluding lower types, the designer forces nature, who is bounded to choose distributions with average k, to place more weight to higher types in equilibrium. Put differently, exclusion is the way the designer finds to curb its excessive pessimism regarding the distribution of types it will face. Such distortion in the extensive margin allows for a mechanism that – compared to what would prevail if positive likelihood were placed to low types – entails less distortions at the intensive margin. In the longer version of this work that is being prepared, we show that such exclusion of types – which has an obvious interpretation in screening and

[7] In fact, in Bayesian models, if $\frac{\partial s}{\partial q}(0, \theta)$ is sufficiently large, not types will be excluded.

taxation models – corresponds, in a Laffont and Tirole (1986) type of setting, to the offer of cost-plus contracts to the least efficient firms. This might suggest that, compared to Bayesian settings, a robust mechanism yields less powerful incentives. Again, in the longer version of this work that is being prepared, we make this suggestion precise and confirm that robustness reduces the power of incentives with which an agent is confronted.

2.2 An example

Let $u(q, \theta) = \theta q$ and $c(q) = q^2/2b$, where $c > 0$. The first-best solution is given by $q^{FB}(\theta) = b\theta$. Solving (3) we get

$$\theta^{\xi}(q) = \exp\,(\xi^{-1}q) \left[(b\xi)^{-1} \int_q^b \exp\,(-\xi^{-1}x)x dx + \exp\,(-\xi^{-1}b) \right]$$

for each $\xi > 0$ and $q \in [0, b]$. Integrating by parts twice, we get

$$\theta^{\xi}(q) = b^{-1} \left(q + \xi \left[1 - \exp\,(\xi^{-1}(q - b)) \right] \right).$$

Hence,

$$\tilde{\theta}(\xi) = \theta^{\xi}(0) = b^{-1}\xi \left[1 - \exp\,\left(-\xi^{-1}b\right) \right].$$

Notice that $\lim\limits_{\xi \to 0} \tilde{\theta}(\xi) = 0$ and $\lim\limits_{\xi \to \infty} \tilde{\theta}(\xi) = 1$. In order to find the multiplier $\xi^* > 0$ we must solve the following equation:

$$\int_0^b \exp\,(\xi^{*-1}q) \left[1 - \exp\,(\xi^{*-1}(q - b)) \right] dq = kb.$$

3 Conclusion

We have considered a general mechanism design problem in which a principal only knows the first moment of the distribution of types of the agent and, in face of such uncertainty, maximizes her worst case expected payoffs. The results and their interpretations have been discussed in the introduction and the main text, so we conclude with extensions that we are currently working on.

As written in the note before the Introduction, two extensions are being prepared. First, we are working on a dynamic version of the main model in which the agent's information evolves over time and all is known by the principal is that types satisfy a martingale condition. We are able to prove that the optimal dynamic robust mechanism is the period by period repetition of the static mechanism this note derived. Second, we are working on a multidimensional, additive, version of the current model. We are able to combine the results of Carroll (2015b) and those derived in this paper to show that the optimal robust multidimensional mechanism entails full separability. Last, we are working out regulation and taxation applications of the results we have established here.

References

Bergemann, D. and K. Schlag (2008): "Pricing without priors," *Journal of the European Economic Association*, 6, 560–569.

——— (2011): "Robust monopoly pricing," *Journal of Economic Theory*, 146, 2527–2543.

Bose, S., E. Ozdenoren, and A. Pape (2006): "Optimal auctions with ambiguity," *Theoretical Economics*, 1, 411–438.

Caldentey, R., Y. Liu, and I. Lobel (2015): "Intertemporal pricing under minimax regret," *mimeo, NYU Stern.*

Carrasco, V., V. Luz, P. Monteiro, and H. Moreira (2015): "Robust Selling Mechanism," *mimeo.*

Carrasco, V. and H. Moreira (2013): "Robust decision-making," *mimeo, PUC-Rio.*

Carroll, G. (2013): "Notes on informationally robust monopoly pricing," *mimeo, Stanford University.*

——— (2015a): "Robustness and linear contracts," *American Economic Review*, 105, 536–563.

——— (2015b): "Robustness and Separation in Multidimensional Screening," *mimeo, Stanford University.*

Frankel, A. (2014): "Aligned delegation," *American Economic Review*, 104, 66–83.

Garrett, D. (2014): "Robustness of simple menus of contracts in cost-based procurement," *Games and Economic Behavior*, 87, 631–641.

Handel, B. and K. Misra (forthcoming): "Robust new product pricing," *Marketing Science.*

Hurwicz, L. and L. Shapiro (1978): "Incentive structures maximizing residual gain under incomplete information," *The Bell Journal of Economics*, 9, 180–191.

Laffont, J. and J. Tirole (1986): "Using cost observation to regulate firms," *Journal of Political Economy*, 3, 614–641.

Mirrlees, J. (1971): "An exploration in the theory of optimum income taxation," *Review of Economic Studies*, 38, 175–208.

Mussa, M. and S. Rosen (1978): "Monopoly and product quality," *Journal of Economic Theory*, 18, 301–317.

Wolitsky, A. (2014): "Mechanism design with maxmin agents: theory and an application to bilateral trade," *mimeo, MIT*.

Departamento de Economia PUC-Rio
Pontifícia Universidade Católica do Rio de Janeiro
Rua Marques de Sâo Vicente 225 - Rio de Janeiro 22453-900, RJ
Tel.(21) 31141078 Fax (21) 31141084
www.econ.puc-rio.br
flavia@econ.puc-rio.br