

# Bayesian incentive compatible parametrization of mechanisms

Thomas A. Weber<sup>a,\*</sup>, Abhishek Bapna<sup>b</sup>

<sup>a</sup> *Department of Management Science and Engineering, Terman Engineering Center,  
Stanford University, Stanford, CA 94305-4026, United States*

<sup>b</sup> *Google, Inc., 1600 Amphitheatre Parkway, Mountain View, CA 94043, United States*

Received 29 April 2006; received in revised form 25 October 2007; accepted 27 October 2007

Available online 19 November 2007

---

## Abstract

We consider a general scheme to construct Bayesian incentive-compatible mechanisms using a suitable ‘variable mechanism parametrization.’ The key idea is to perturb a given direct mechanism, which might not be truth revealing, introducing sufficient variability as a function of agents’ announcements to generate incentives for truthful revelation. We discuss a variable-price auction in a general setting as an example.

© 2007 Elsevier B.V. All rights reserved.

*JEL Classification:* D42; D44; D82; D83

*Keywords:* Bayesian incentive compatibility; Direct revelation mechanisms; Mechanism design

---

Let the will be set on the Way.  
Confucius (K’ung fu-tzu) (551–479 BC).<sup>1</sup>

## 1. Introduction

Multi-agent mechanism design has proved useful for the implementation of allocation rules in many diverse settings, such as public decision making, multilateral trade, or the distribution of resources within a firm. A standard approach to mechanism design uses the revelation principle (Gibbard, 1973; Myerson, 1979), which allows limiting attention to direct mechanisms,<sup>2</sup> subject to Bayesian incentive-compatibility constraints which guarantee that truthfully announcing private type information to the principal (who then makes a resource-allocation decision based on all announcements received) is an undominated strategy for all agents participating in the mechanism. Real-world mechanisms, on the other hand, typically are neither direct nor do they induce truthful revelation of private information. This may be due to a lack of design (i.e., bounded rationality), to institutional limitations (which hamper the consideration of certain agent heterogeneities), or, perhaps, to a conscious effort by the mechanism designer to strongly implement a particular

---

\* Corresponding author. Tel.: +1 650 725 6827.

E-mail addresses: [webert@stanford.edu](mailto:webert@stanford.edu) (T.A. Weber), [abapna@google.com](mailto:abapna@google.com) (A. Bapna).

<sup>1</sup> Analects 7.6.

<sup>2</sup> A *direct* mechanism is such that an agent’s possible announcements must come from the set that contains this agent’s private type information (i.e., his type space).

social welfare function, circumventing certain undesirable equilibria introduced by applying the revelation principle (Palfrey, 2002). Whatever the reason may be, we start from the premise that existing mechanisms are at least somewhat flexible and contain parameters that can be adjusted. The question we then address is *how to systematically use the available flexibility in existing mechanisms in order to achieve Bayesian incentive compatibility*. The central idea is to use an announcement-contingent parameter vector, which may be different for each agent (depending on observable heterogeneities). The parameter-induced systematic change in the agents' payoffs as a function of their announcements may be used to provide the missing incentives for truthful revelation of their private information. This approach is rather general: for instance, when using the parameters to establish a homotopy relation between certain 'extreme' mechanisms (which provide 'extreme' incentives), it is often possible to interpret our method as an announcement-contingent *randomization* of mechanisms,<sup>3</sup> which may prove exceedingly useful in complex environments, e.g., when designing efficient dynamic allocation mechanisms (Bapna and Weber, 2005).<sup>4</sup> A key notion we use to characterize 'extreme' incentives is one-sided (upward/downward) Bayesian incentive compatibility, expressing the fact that agents *either* never under- *or* never over-report their types to the principal under a given mechanism. Providing a general framework to combine such one-sided mechanisms, or, more generally, perturbing a non-incentive-compatible mechanism to yield truthful mechanisms is the essential contribution of this paper. We illustrate our approach by constructing a truthful variable-price auction in a general asymmetric common-values setting.

## 2. The model

### 2.1. Preliminaries

Suppose there are  $n \geq 2$  agents and one mechanism-designing principal. An agent  $i \in \mathcal{N} = \{1, \dots, n\}$  is of type  $\theta^i \in \Theta^i = [\underline{\theta}^i, \bar{\theta}^i]$ , where  $\underline{\theta}^i$  and  $\bar{\theta}^i$  are finite real numbers with  $\underline{\theta}^i < \bar{\theta}^i$ . The agents' types are jointly distributed on  $\Theta$ . At the start of the game each agent  $i$  privately observes his own type  $\theta^i$ , which determines his utility  $U^i((\xi, \tau); \theta^i)$  for any  $l$ -bundle of commodities  $\xi \in \mathbb{R}_+^l$  (where  $l \geq 1$  is an integer) and any monetary transfer  $\tau \in \mathbb{R}$  to the principal. We assume that for any  $i \in \mathcal{N}$  the utility function  $U^i(\cdot; \theta^i) : \mathbb{R}_+^l \times \mathbb{R} \rightarrow \mathbb{R}$  is smooth on  $\mathbb{R}_+^l \times \mathbb{R}$  and is smoothly parameterized on  $\Theta^i$ .

Let  $\alpha = (\alpha^1, \dots, \alpha^n)$  be an allocation rule, where for each  $i \in \mathcal{N}$  the function  $\alpha^i = (x^i, t^i)$  is such that  $x^i : \Theta \rightarrow \mathbb{R}^l$  maps a type vector  $\theta = (\theta^1, \dots, \theta^n)$  in  $\Theta$  to an  $l$ -bundle of commodities  $x^i(\theta)$  in  $\mathbb{R}_+^l$ , while  $t^i : \Theta \rightarrow \mathbb{R}$  maps  $\theta$  to a real-valued transfer payment  $t^i(\theta)$  from agent  $i$  to the principal.<sup>5</sup> We assume that the principal is able to commit to an allocation rule  $\alpha$  which is known to all agents, before each agent  $i$  makes a type announcement  $\hat{\theta}^i \in \Theta^i$  to the principal, who then ensures that the allocation  $\alpha(\hat{\theta})$  with  $\hat{\theta} = (\hat{\theta}^1, \dots, \hat{\theta}^n)$  takes place. This sequential allocation game corresponds to a direct mechanism  $\mathfrak{M} = \langle \alpha, \Theta \rangle$  in which the space of possible agent announcements coincides with the type space  $\Theta$ . The revelation principle ensures that restricting attention to such mechanisms is without loss of generality, as long as the allocation function is *Bayesian incentive compatible* (BIC) in the sense that

$$\forall \theta^i \in \Theta^i : \theta^i \in \arg \max_{\hat{\theta}^i \in \Theta^i} \bar{U}^i(\hat{\theta}^i, \theta^i), \quad (1)$$

where

$$\bar{U}^i(\hat{\theta}^i, \theta^i) = E_{\theta_{-i}}[U^i(\alpha^i(\hat{\theta}^i, \bar{\theta}^{-i}); \theta^i) | \theta^i] \quad (2)$$

denotes agent  $i$ 's expected utility conditional on truthful announcements by other agents (given his beliefs about their random types  $\bar{\theta}^{-i}$ ), for all  $\hat{\theta}^i, \theta^i \in \Theta^i$ . While agent  $i$ 's expected utility  $\bar{U}^i$  clearly depends on the allocation rule  $\alpha^i$ , we choose to not explicitly acknowledge this dependency on the left-hand side of (2). The reasons for this

<sup>3</sup> The idea of randomized mechanisms is not new and dates back at least to Myerson (1981).

<sup>4</sup> We limit attention to Bayesian implementation (Postlewaite and Schmeidler, 1986; Palfrey and Srivastava, 1993). When agents' utilities are quasilinear, dominant-strategy incentive compatibility can be implemented without budget balance through Vickrey–Clarke–Groves mechanisms (Vickrey, 1961; Clarke, 1971; Groves, 1973), and Bayesian incentive compatibility can be implemented with budget balance through d'Aspremont–Gérard-Varet–Arrow (AGV–Arrow) mechanisms (Arrow, 1979; d'Aspremont and Gérard-Varet, 1979). In the absence of quasilinearity this is generically not possible, regardless of budget balance.

<sup>5</sup> We introduce the principal to ensure budget balance, even when utilities are not quasilinear (so that the AGV–Arrow mechanism offers no remedy).

are that (i) the function  $\alpha^i$  remains fixed throughout the paper, and that (ii) we like to emphasize an announcement-dependent reparametrization of the original allocation rule which is introduced as follows: consider a direct mechanism  $\mathfrak{M}(\bar{\eta}) = (\alpha^1(\cdot; \bar{\eta}^1), \dots, \alpha^n(\cdot; \bar{\eta}^n), \Theta)$  whose allocation function depends on the value of a parameter matrix  $\bar{\eta} = [\bar{\eta}_k^i]_{(i,k) \in \mathcal{N} \times \mathcal{M}}$ , where  $\mathcal{M} = \{1, \dots, m\}$  and  $m \geq 2$  is an integer that denotes the number of parameters per agent.<sup>6</sup> Then, for any given  $\bar{\eta} = (\bar{\eta}_1^i, \dots, \bar{\eta}_m^i)$ , agent  $i$ 's set of optimal announcements is

$$\mathcal{A}^i(\theta^i; \bar{\eta}^i) = \arg \max_{\hat{\theta}^i \in \Theta^i} \bar{U}^i(\hat{\theta}^i, \theta^i; \bar{\eta}^i),$$

so that Bayesian incentive compatibility is equivalent to the inclusion  $\theta^i \in \mathcal{A}^i(\theta^i; \bar{\eta}^i)$  for all  $\theta^i \in \Theta^i$  and all  $i \in \mathcal{N}$ .

Assume now that the principal, instead of fixing a parameter matrix  $\bar{\eta}$  at the outset, makes the  $i$ th row vector which we term  $\eta^i = (\eta_1^i, \dots, \eta_m^i)$  of a matrix  $\eta = [\eta_k^i]_{(i,k) \in \mathcal{N} \times \mathcal{M}}$  dependent on agent  $i$ 's announcement, i.e.,

$$\eta^i = \eta^i(\hat{\theta}^i).$$

This includes the special case of a *fixed* parametrization with  $\eta^i(\hat{\theta}^i) \equiv \bar{\eta}^i$  discussed above.

**Definition 1** (*Fixed/variable parametrization*). For each  $i \in \mathcal{N}$  let  $E^i \subset \mathbb{R}^m$  be a locally path connected topological space and let  $E = E^1 \times \dots \times E^n$ . The mechanism parametrization  $\mathfrak{M}(\eta)$  with  $\eta = (\eta^1, \dots, \eta^n)$  is called a *fixed* mechanism parametrization (on  $E$ ) if the mapping  $\eta^i : \Theta^i \rightarrow E^i$  is constant with  $\eta^i(\hat{\theta}^i) \equiv \bar{\eta}^i \in \mathbb{R}^m$  for all  $i \in \mathcal{N}$ . Otherwise it is called a *variable* mechanism parametrization (on  $E$ ).

Thus, instead of the parameterized mechanism  $\mathfrak{M}(\bar{\eta})$  for a fixed parameter, the principal offers the mechanism  $\mathfrak{M}(\eta)$  for a variable parameter, so that Bayesian incentive compatibility amounts to

$$\theta^i \in \mathcal{A}^i(\theta^i; \eta^i) = \arg \max_{\hat{\theta}^i \in \Theta^i} \bar{U}^i(\hat{\theta}^i, \theta^i; \eta^i(\hat{\theta}^i)) \tag{3}$$

for all  $\theta^i \in \Theta^i$  and all  $i \in \mathcal{N}$ .

### 2.2. Statement of the problem

Using the terminology in Definition 1, we are interested in conditions under which for a given fixed mechanism parametrization  $\mathfrak{M}(\bar{\eta})$  on  $E$  that does not need to be BIC, a variable mechanism parametrization  $\mathfrak{M}(\eta)$ , containing the mapping  $\eta : \Theta \rightarrow E$ , can be found that is BIC. We refer to this question as problem (P). To deal with the agents' participation constraint (requiring that  $\bar{U}^i(\theta^i, \theta^i; \bar{\eta}^i) \geq \bar{U}_0^i(\theta^i)$  for some smooth expected reservation utility  $\bar{U}_0^i : \Theta^i \rightarrow \mathbb{R}$ ), we assume in what follows that the parameter domain  $E \subset \mathbb{R}^m$  is (i) nonempty, (ii) locally path connected, and (iii) *individually rational with respect to*  $(U^1, \dots, U^n)$  in the sense that

$$\bar{\eta} \in E \Rightarrow (\bar{U}^1(\theta^1, \theta^1; \bar{\eta}^1) - \bar{U}_0^1(\theta^1), \dots, \bar{U}^n(\theta^n, \theta^n; \bar{\eta}^n) - \bar{U}_0^n(\theta^n)) \geq 0 \tag{4}$$

for all  $\theta \in \Theta$ .

### 2.3. Bayesian incentive compatibility

To address problem (P) it is useful to first relax the concept of Bayesian incentive compatibility and express it for any agent  $i \in \mathcal{N}$  with respect to a single-valued measurable selection  $\sigma^i(\theta^i) \in \mathcal{A}^i(\theta^i)$ .

**Definition 2** (*One-sided Bayesian incentive compatibility*). Let  $\mathfrak{M}$  be a direct mechanism. (i) The mechanism  $\mathfrak{M}$  is downward (upward) Bayesian incentive compatible, in short d-BIC (u-BIC), with respect to a selection  $\sigma = (\sigma^1, \dots, \sigma^n)$  if for any  $i \in \mathcal{N}$  and any  $\theta^i \in \Theta^i$  we have that

$$\sigma^i(\theta^i) \geq (\leq) \theta^i. \tag{5}$$

<sup>6</sup> The requirement that  $m \geq 2$  is without loss of generality, since the dependence of the mechanism on  $\eta_k^i$  for  $k \geq 2$  can be trivial for any  $i \in \mathcal{N}$ .

(ii) The mechanism  $\mathfrak{M}$  is downward (upward) incentive compatible if it is d-BIC (u-BIC) with respect to any selection  $\sigma \in \mathcal{A}$ . (iii) The mechanism  $\mathfrak{M}$  is one-sided if it is either d-BIC or u-BIC.

**Lemma 1 (BIC).** *A direct mechanism  $\mathfrak{M}$  is BIC if and only if there exists a selection  $\sigma \in \mathcal{A}$  such that  $\mathfrak{M}$  is both d-BIC and u-BIC with respect to  $\sigma$ .*

**Proof.** BIC implies by (1) both d-BIC and u-BIC. The converse is true, since  $\theta^i \leq \sigma^i(\theta^i) \leq \theta^i$  implies that  $\sigma^i(\theta^i) = \theta^i \in \mathcal{A}^i(\theta^i)$  for all  $i \in \mathcal{N}$  and all  $\theta^i \in \Theta^i$ .  $\square$

We provide a characterization of Bayesian incentive compatibility which is adapted to our context and then, using an approach similar to the one in Weber (2005), derive a ‘sufficient BIC criterion’ which proves useful in finding a variable mechanism parametrization that is Bayesian incentive compatible.

**Lemma 2 (Equivalent BIC criterion).** *The parameterized mechanism  $\mathfrak{M}(\eta)$  is BIC if and only if<sup>7</sup>*

$$\frac{\partial \bar{U}^i}{\partial \hat{\theta}^i} \Big|_{(\theta^i, \theta^i; \eta^i(\theta^i))} + \frac{\partial \bar{U}^i}{\partial \eta^i} \Big|_{(\theta^i, \theta^i; \eta^i(\theta^i))} \cdot \dot{\eta}^i(\theta^i) = 0 \tag{6}$$

and

$$(\hat{\theta}^i - \theta^i) \left( \frac{\partial \bar{U}^i}{\partial \theta^i} \Big|_{(\hat{\theta}^i, \theta^i; \eta^i(\hat{\theta}^i))} - \frac{\partial \bar{U}^i}{\partial \theta^i} \Big|_{(\theta^i, \theta^i; \eta^i(\theta^i))} \right) \geq 0. \tag{7}$$

**Proof.**  $\Rightarrow$ : If  $\mathfrak{M}(\eta)$  is BIC, then relation (3) implies, by Fermat’s rule, condition (6) for all  $\theta \in \text{int}\Theta$ , so that, by continuity of the left-hand side, condition (6) is satisfied for all  $\theta \in \Theta$ . This implies that the total derivative

$$\frac{d\bar{U}^i(\theta^i, \theta^i; \eta^i(\theta^i))}{d\theta^i} = \frac{\partial \bar{U}^i}{\partial \theta^i} \Big|_{(\theta^i, \theta^i; \eta^i(\theta^i))},$$

whence we obtain, by the fundamental theorem of calculus (Rudin, 1976, p. 134), that

$$\bar{U}^i(\hat{\theta}^i, \hat{\theta}^i; \eta^i(\hat{\theta}^i)) - \bar{U}^i(\theta^i, \theta^i; \eta^i(\theta^i)) = \int_{\theta^i}^{\hat{\theta}^i} \frac{\partial \bar{U}^i}{\partial \theta^i} \Big|_{(\vartheta, \vartheta; \eta^i(\vartheta))} d\vartheta$$

for all  $\hat{\theta}^i, \theta^i \in \Theta^i$  and all  $i \in \mathcal{N}$ . Again, by the fundamental theorem of calculus, we have that

$$\bar{U}^i(\hat{\theta}^i, \hat{\theta}^i; \eta^i(\hat{\theta}^i)) - \bar{U}^i(\hat{\theta}^i, \theta^i; \eta^i(\hat{\theta}^i)) = \int_{\theta^i}^{\hat{\theta}^i} \frac{\partial \bar{U}^i}{\partial \theta^i} \Big|_{(\hat{\theta}^i, \vartheta; \eta^i(\vartheta))} d\vartheta$$

for all  $\hat{\theta}^i, \theta^i \in \Theta^i$  and all  $i \in \mathcal{N}$ . Since BIC of  $\mathfrak{M}(\eta)$  is equivalent to

$$\bar{U}^i(\theta^i, \theta^i; \eta^i(\theta^i)) \geq \bar{U}^i(\hat{\theta}^i, \theta^i; \eta^i(\hat{\theta}^i)), \tag{8}$$

we find that

$$\int_{\theta^i}^{\hat{\theta}^i} \left[ \frac{\partial \bar{U}^i}{\partial \theta^i} \Big|_{(\hat{\theta}^i, \vartheta; \eta^i(\vartheta))} - \frac{\partial \bar{U}^i}{\partial \theta^i} \Big|_{(\vartheta, \vartheta; \eta^i(\vartheta))} \right] d\vartheta \geq 0$$

for all  $\hat{\theta}^i, \theta^i \in \Theta^i$  and all  $i \in \mathcal{N}$ . The last inequality holds if and only if condition (7) is satisfied.  $\Leftarrow$ : From our earlier arguments we obtain directly that conditions (6) and (7) together imply (8), i.e., the Bayesian incentive compatibility of  $\mathfrak{M}(\eta)$ .  $\square$

<sup>7</sup> For notational convenience, we use dots to indicate total derivatives.

**Corollary 1** (Sufficient BIC criterion). *The parameterized mechanism  $\mathfrak{M}(\eta)$  is BIC if for every  $i \in \mathcal{N}$  condition (6) holds on  $\Theta^i$  and*

$$G^i(\hat{\theta}^i, \eta^i(\hat{\theta}^i), \dot{\eta}^i(\hat{\theta}^i)) \geq 0 \quad (9)$$

for all  $\hat{\theta}^i \in \Theta^i$ , where

$$G^i(\hat{\theta}^i, \bar{\eta}^i, v) = \min_{\theta^i \in \Theta^i} \left\{ \frac{\partial^2 \bar{U}^i}{\partial \hat{\theta}^i \partial \theta^i} \Big|_{(\hat{\theta}^i, \theta^i, \bar{\eta}^i)} + \frac{\partial^2 \bar{U}^i}{\partial \theta^i \partial \eta^i} \Big|_{(\hat{\theta}^i, \theta^i, \bar{\eta}^i)} \cdot v \right\} \quad (10)$$

for all  $(\hat{\theta}^i, \bar{\eta}^i, v) \in \Theta^i \times E^i \times \mathbb{R}^m$ .

Note that the single-crossing condition (7) is not particularly useful in practice, so that we limit attention to the sufficient condition (9) which implies (7). If the agents' expected utilities are supermodular in announcements and types, this is also without any loss in generality.

### 3. Main result

Let us first introduce the control variable  $u^i = (u_1^i, \dots, u_m^i) = (\dot{\eta}_1^i, \dots, \dot{\eta}_m^i)$ . To state our results in a concise way it is convenient to assume that  $\bar{U}^i(\hat{\theta}^i, \theta^i; \eta^i)$  is strictly increasing in  $\eta_1^i$ .<sup>8</sup> We can then rewrite the sufficient implementability conditions (6) and (9) from Corollary 1 in the form

$$\dot{\eta}^i(\theta^i) = \varphi^i(\theta^i, \eta^i(\theta^i), u_{-1}^i), \quad (11)$$

$$u_{-1}^i(\theta^i) \in \mathcal{U}_{-1}^i(\theta^i, \eta^i(\theta^i)), \quad (12)$$

for all  $\theta^i \in \Theta^i$  and  $i \in \mathcal{N}$ , where  $u_{-1}^i = (u_2^i, \dots, u_m^i)$ ,

$$\varphi^i(\theta^i, \bar{\eta}^i, u_{-1}^i) = \begin{bmatrix} - \left( \frac{\partial \bar{U}^i}{\partial \eta_1^i} \Big|_{(\theta^i, \theta^i, \bar{\eta}^i)} \right)^{-1} \left( \frac{\partial \bar{U}^i}{\partial \hat{\theta}^i} \Big|_{(\theta^i, \theta^i, \bar{\eta}^i)} + \sum_{k=2}^m \frac{\partial \bar{U}^i}{\partial \eta_k^i} \Big|_{(\theta^i, \theta^i, \bar{\eta}^i)} u_k^i \right) \\ u_2^i \\ \vdots \\ u_m^i \end{bmatrix}, \quad (13)$$

and

$$\mathcal{U}_{-1}^i(\theta^i, \bar{\eta}^i) = \{\hat{u}_{-1}^i \in \mathbb{R}^{m-1} : G^i(\theta^i, \bar{\eta}^i, \varphi^i(\theta^i, \bar{\eta}^i, \hat{u}_{-1}^i)) \geq 0\} \quad (14)$$

for all  $(\theta^i, \bar{\eta}^i, u_{-1}^i) \in \Theta^i \times E^i \times \mathbb{R}^{m-1}$ .

**Theorem 1.** *Let  $\bar{\eta}_d, \bar{\eta}_u \in E$ . The variable mechanism parametrization  $\mathfrak{M}(\eta)$  is BIC if the following three conditions are satisfied: (C1) the fixed mechanism parametrization  $\mathfrak{M}(\bar{\eta}_d)$  is  $d$ -BIC; (C2) the fixed mechanism parametrization  $\mathfrak{M}(\bar{\eta}_u)$  is  $u$ -BIC; and (C3) for any  $i \in \mathcal{N}$  the system (11) can be steered from  $\eta^i(\underline{\theta}^i) = \bar{\eta}_u^i$  to  $\eta^i(\bar{\theta}^i) = \bar{\eta}_d^i$  using  $u_{-1}^i(\theta^i) \in \mathcal{U}_{-1}^i(\theta^i, \eta^i(\theta^i))$  such that  $\eta^i(\theta^i) \in E^i$  for all  $\theta^i \in \Theta^i$ .*

**Proof.** By Definition 2, for any selection  $\sigma^i(\theta^i) \in \mathcal{A}^i(\theta^i; \bar{\eta}_d^i)$  we have that agent  $i$ 's optimal announcement  $\hat{\theta}^i = \sigma^i(\theta^i)$  is not less than his type  $\theta^i$  for almost all  $\theta^i \in \Theta^i$ . Hence, at the upper end of the type interval  $\Theta^i$  it is  $\sigma^i(\bar{\theta}^i) = \bar{\theta}^i$ , i.e., any agent  $i$  of type  $\bar{\theta}^i$  announces truthfully under the mechanism  $\mathfrak{M}(\bar{\eta}_d)$ . Similarly, any agent  $i$  of type  $\underline{\theta}^i$  announces truthfully under the mechanism  $\mathfrak{M}(\bar{\eta}_u)$ . The differential Eq. (11) is equivalent to condition (6) (since  $(\dot{\eta}_1^i, \dots, \dot{\eta}_m^i) = (u_1^i, \dots, u_m^i)$ ), and the inclusion  $\dot{\eta}^i(\theta^i) \in \mathcal{U}_{-1}^i(\theta^i, \eta^i(\theta^i))$  is equivalent to condition (7).  $\square$

<sup>8</sup> All we need is that relation (6),  $\partial \bar{U}^i / \partial \hat{\theta}^i + \partial \bar{U}^i / \partial \eta^i \cdot (u_1^i, \dots, u_m^i) = 0$ , can be solved for  $u_1^i$  (or, without loss of generality, for any other single component of  $u^i$ ) almost everywhere on  $\Theta^i$ .

**Theorem 1** requires that the state  $\eta^i(\theta^i)$  of the system described by (11) can be steered from  $\eta^i(\underline{\theta}^i) = \bar{\eta}_d^i$  to  $\eta^i(\bar{\theta}^i) = \bar{\eta}_u^i$ , subject to the constraint (12) and the state constraint  $\eta^i(\theta^i) \in E^i$ . Consider now agent  $i$ . If any other agent  $j$  follows a truthful announcement strategy, then it is clear that *either* starting *or* ending the trajectory at a one-sided mechanism, subject to the last two constraints, is enough to guarantee the Bayesian incentive compatibility of the parametrization  $\mathfrak{M}(\eta)$ . We can thus strengthen **Theorem 1**.

**Theorem 1’.** *Let  $\bar{\eta}_d, \bar{\eta}_u \in E$ . The variable mechanism parametrization  $\mathfrak{M}(\eta)$  is BIC if, in addition to (C3), either (C1) or (C2) in **Theorem 1** is satisfied.*

The last theorem states that if a parameter path  $\eta(\theta)$  either starts or ends at a point  $\bar{\eta} \in E$  at which the mechanism  $\mathfrak{M}(\bar{\eta})$  is one-sided, then the variable mechanism parametrization  $\mathfrak{M}(\eta)$  is BIC. One-sidedness at an endpoint is useful in guaranteeing that all agents find it optimal there to report their types truthfully.

**4. Application: truthful variable-price auctions**

To illustrate our findings, let us consider a positive linear combination of  $m$  different  $k$ -price auctions (where  $k \in \{1, \dots, m\}$  and  $m \geq 3$ ) in a general asymmetric common-values setting. Suppose that there are  $n \geq m$  risk-averse agents, jointly distributed on the full support  $\Theta = \Theta^1 \times \dots \times \Theta^n$  with smooth joint density  $f$ . To obtain convenient stochastic ordering properties, we make the standard assumption that types are affiliated (Milgrom and Weber, 1982; Karlin and Rinott, 1980), i.e., that  $f$  is log-supermodular. Consistent with (weak) risk aversion, we assume that for any realization  $\theta = (\theta^1, \dots, \theta^n)$  of the random vector  $\tilde{\theta} = (\tilde{\theta}^1, \dots, \tilde{\theta}^n)$ , any agent  $i$ ’s utility  $U^i(\cdot; \theta^i) : \mathbb{R} \rightarrow \mathbb{R}$ , where  $\theta^i \in \Theta^i = [\underline{\theta}^i, \bar{\theta}^i] \subset \mathbb{R}_{++}$ , is a smooth concave function. The mechanism is parameterized such that the weight  $\eta_k^i$  on the  $k$ th highest price for agent  $i$  is a mapping from  $\Theta^i$  to nonnegative real numbers, so that any coefficient vector  $\eta^i = (\eta_1^i, \dots, \eta_m^i)$  is an element of

$$E^i = \{\bar{\eta}^i \in \mathbb{R}_+^m : \min_{\theta^i \in \Theta^i} \bar{U}^i(\theta^i, \theta^i, \bar{\eta}^i) \geq 0\},$$

which is nonempty, locally path connected, and individually rational with respect to  $U^i$ , assuming that all agents’ reservation utility is zero.

We are now ready to construct a variable mechanism parametrization of this mechanism that is BIC. If all agents other than agent  $i$  follow a truthful bidding strategy (so that  $\hat{\theta}^j \equiv b(\theta^j) \equiv \theta^j$  for all  $j \neq i$ ), agent  $i$ ’s expected utility from announcing  $\hat{\theta}^i \in \Theta^i$  is<sup>9</sup>

$$\bar{U}^i(\hat{\theta}^i, \theta^i; \eta^i) = P(\tilde{\theta}^{(1)} = \hat{\theta}^i) E[U^i(\rho^i(\hat{\theta}^i, \theta^i, \tilde{\theta}; \eta^i); \theta^i) | \tilde{\theta}^{(1)} = \hat{\theta}^i, \tilde{\theta}^i = \theta^i],$$

where

$$\rho^i(\hat{\theta}^i, \theta^i, \tilde{\theta}; \eta^i) = \theta^i - \eta_1^i(\hat{\theta}^i)\hat{\theta}^i - \sum_{k=2}^m \eta_k^i(\hat{\theta}^i)b(\tilde{\theta}^{(k)}).$$

For any  $i \in \mathcal{N}$ , the implementability condition (6) then becomes

$$\int_{\Theta^{-i}} \left( \eta_1^i(\theta^i) + \sum_{k=1}^m \tilde{\theta}^{(k)} \eta_k^i(\theta^i) \right) U^i(\rho^i(\theta^i, \theta^i, \tilde{\theta}; \eta^i); \theta^i) f(\tilde{\theta} | \tilde{\theta}^{(1)} = \tilde{\theta}^i = \theta^i) d\tilde{\theta}^{-i} = R^i(\theta^i; \eta^i),$$

where

$$R^i(\theta^i; \eta^i) = \frac{\bar{U}^i(\theta^i, \theta^i; \eta^i)}{P(\tilde{\theta}^{(1)} = \theta^i)} \left( \frac{\partial \ln P(\tilde{\theta}^{(1)} = \theta^i)}{\partial \theta^i} \right) + \int_{\Theta^{-i}} U^i(\rho^i(\theta^i, \theta^i, \tilde{\theta}; \eta^i); \theta^i) \frac{\partial f(\tilde{\theta} | \tilde{\theta}^{(1)} = \hat{\theta}^i, \tilde{\theta}^i = \theta^i)}{\partial \hat{\theta}^i} \Bigg|_{\hat{\theta}^i = \theta^i} d\tilde{\theta}^{-i}$$

is the change of agent  $i$ ’s expected utility in case of winning due to the effect a marginal increase of his announcement has on the probability of winning. In equilibrium, the latter equals the marginal change of agent  $i$ ’s payment to the

<sup>9</sup> We denote by  $a^{(k)}$  the  $k$ th largest element of a given vector  $a$ .

principal. The last condition, together with inequality (9), is by Corollary 1 sufficient for BIC. Note that in our setting for any  $\bar{\eta}^i \in E^i$  and  $v = (v_1, \dots, v_m) \in \mathbb{R}^m$  we have that

$$G^i(\hat{\theta}^i, \bar{\eta}^i, v) = P(\tilde{\theta}^{(1)} = \hat{\theta}^i) \min_{\theta^i \in \Theta^i} \left\{ E \left[ \left( \frac{\partial \ln P(\tilde{\theta}^{(1)} = \hat{\theta}^i)}{\partial \hat{\theta}^i} \right) \dot{U}^i(\rho^i; \theta^i) - \bar{\eta}_1^i \dot{U}^i(\rho^i; \theta^i) \middle| \tilde{\theta}^{(1)} = \hat{\theta}^i, \tilde{\theta}^i = \theta^i \right] \right. \\ \left. + \int_{\Theta^{-i}} \dot{U}^i(\rho^i; \theta^i) \frac{\partial f(\tilde{\theta} | \tilde{\theta}^{(1)} = \hat{\theta}^i, \tilde{\theta}^i = \theta^i)}{\partial \hat{\theta}^i} \middle|_{\hat{\theta}^i = \theta^i} d\tilde{\theta}^{-i} \right. \\ \left. - \sum_{k=1}^m v_k E \left[ \tilde{\theta}^{(k)} \dot{U}^i(\rho^i; \theta^i) \middle| \tilde{\theta}^{(1)} = \hat{\theta}^i, \tilde{\theta}^i = \theta^i \right] \right\}$$

where  $\rho^i = \rho^i(\hat{\theta}^i, \theta^i, \tilde{\theta}; \bar{\eta}^i)$ . The evolution of  $\eta^i$  in the variable mechanism parametrization is governed by relations (11) and (12), where

$$\varphi^i(\theta^i, \bar{\eta}^i, u_{-1}^i) = \begin{bmatrix} \frac{\bar{\eta}_1^i}{\theta^i} - \frac{R^i(\theta^i; \bar{\eta}^i) - \sum_{k=2}^m u_k^i E[\tilde{\theta}^{(k)} \dot{U}^i(\rho^i(\theta^i, \theta^i, \tilde{\theta}; \bar{\eta}^i); \theta^i) | \tilde{\theta}^{(1)} = \tilde{\theta}^i = \theta^i]}{\theta^i E[\dot{U}^i(\rho^i(\theta^i, \theta^i, \tilde{\theta}; \bar{\eta}^i); \theta^i) | \tilde{\theta}^{(1)} = \tilde{\theta}^i = \theta^i]} \\ u_2^i \\ \vdots \\ u_m^i \end{bmatrix}. \tag{15}$$

If we set  $\bar{\eta}_u^i = (1, 0, \dots, 0)$  and let  $\bar{\eta}_d^i \in E^i$  be free, then the fixed mechanism parametrization  $\mathfrak{M}(\bar{\eta}_u^i)$  is u-BIC (since a first-price auction is u-BIC, resulting in the familiar ‘bid shaving’). Thus, condition (C2) is satisfied. We now check condition (C3), restricting attention to ‘controls’  $u_{-1}^i \in \mathbb{R}^{m-1}$  such that  $u_2 = \dots = u_{m-1} = 0$ . Since  $f$  is log-supermodular by assumption, the sum of the first two terms of the minimand in the definition of  $G^i$  is nonnegative. Since  $E[(\tilde{\theta}^{(1)} - \tilde{\theta}^{(m)}) \dot{U}^i(\rho^i(\theta^i, \theta^i, \tilde{\theta}; \bar{\eta}^i); \theta^i) | \tilde{\theta}^{(1)} = \tilde{\theta}^i = \theta^i] \geq 0$ , the last term is nonnegative, as long as  $v_m \geq v_1$ . But the previous inequality is satisfied if  $u_m^i$  is larger than the first entry of  $\varphi^i$  (termed  $\varphi_1^i$ , corresponding to  $u_1^i$ ) in (15), or, equivalently, if given  $\bar{\eta}^i$  and  $\theta^i \in \Theta^i$  it is

$$u_m^i(\theta^i) \geq - \frac{\bar{\eta}_1^i E[\dot{U}^i(\rho^i(\theta^i, \theta^i, \tilde{\theta}; \bar{\eta}^i); \theta^i) | \tilde{\theta}^{(1)} = \tilde{\theta}^i = \theta^i] + R^i(\theta^i; \bar{\eta}^i)}{E[(\tilde{\theta}^{(1)} - \tilde{\theta}^{(m)}) \dot{U}^i(\rho^i(\theta^i, \theta^i, \tilde{\theta}; \bar{\eta}^i); \theta^i) | \tilde{\theta}^{(1)} = \tilde{\theta}^i = \theta^i]} \equiv \underline{u}_m^i(\theta^i, \bar{\eta}^i).$$

Hence, we can conclude using Theorem 1' that, as long as the control-constraint

$$u_m^i(\theta^i) \geq \underline{u}_m^i(\theta^i, \eta^i(\theta^i)) \tag{16}$$

is satisfied on  $\Theta^i$ , the variable mechanism parametrization  $\mathfrak{M}(\eta)$  with

$$\begin{cases} \hat{\eta}^i(\theta^i) = (\varphi_1^i(\theta^i, \eta^i(\theta^i), (0, \dots, u_m^i(\theta^i))), 0, \dots, 0, u_m^i(\theta^i)), \\ \eta^i(\theta^i) = \bar{\eta}_d^i, \end{cases} \tag{17}$$

for all  $i \in \mathcal{N}$  and all  $\theta^i \in \Theta^i$ , is BIC (cf. Fig. 1).<sup>10</sup> Intuitively, by selecting a starting point  $\bar{\eta}_u^i$  and an endpoint  $\bar{\eta}_d^i$  that are in  $E^i$  and making sure that the trajectory never leaves  $E^i$  in between, individual rationality is satisfied by the definition

<sup>10</sup> A solution  $\eta^i$  to the initial value problem (17) subject to  $\eta^i \in E^i$  exists, since we can find a  $u_m^i$  such that the parameter trajectory stays within the





for some  $m \geq 3$  seems rather natural. Note that our results do not depend in any substantive way on the agents' risk preferences (nor on assumptions about the type distribution, except for smoothness). The substantial flexibility available when choosing a parameter path that yields a truthful mechanism (cf. Fig. 1) can be used to achieve surplus distribution objectives that the principal might have; Bapna and Weber (2005) show how this may be accomplished in a specific dynamic mechanism design problem. The main idea for using a BIC mechanism parametrization to shift surplus from the agents to the principal, or between agents, is to first determine an "ideal" parametrization that is individually rational but not necessarily BIC so as to satisfy the revenue distribution objectives (e.g., revenue maximization, efficiency, or fairness). Subsequently, one can approximate the trajectory of the ideal parametrization by a finite number of pieces of BIC trajectories, thereby approaching the desired revenue objective up to a pre-specified error.

## 5. Conclusion

In existing mechanisms parameters often occur naturally and one-sided incentives can easily be induced by providing extreme incentives. To understand what drives the BIC-parametrization idea in this paper it is best to consider **Theorem 1'**, which is a little stronger than **Theorem 1**. It asserts that when starting a mechanism parametrization at a point  $\bar{\eta}$  for which the mechanism  $\mathcal{M}(\bar{\eta})$  is one-sidedly BIC (either u-BIC or d-BIC), then the parametrization can be rendered BIC by steering  $\eta(\theta)$  away from the initial point  $\bar{\eta}$  (as a function of announcements), subject to satisfying the incentive-compatibility constraints (11) and (12). What is important is that one starts with a one-sidedly BIC mechanism, and then stays on a path which guarantees and maintains BIC. The key insight therefore is: in order to achieve Bayesian incentive compatibility, it is mostly the direction of the parametrization that matters. Its absolute position is important only insofar as it ensures that the mechanism along its path stays individually rational. The direction of the path is chosen in terms of the control variable  $u^i = (u_1^i, \dots, u_m^i) = (\dot{\eta}_1^i, \dots, \dot{\eta}_m^i)$  for each agent  $i \in \mathcal{N}$ . As long as the control-constraint set  $\mathcal{U}_{-1}^i(\theta^i, \eta^i(\theta^i))$  is not empty along the chosen path, BIC can be established. Because of this local approach, it is important that it is possible to vary the  $\eta(\theta)$  in a neighborhood by selecting an appropriate slope through the controls. When the type space is discrete, an equivalent BIC criterion cannot be written in the convenient format of **Lemma 2**. Nonetheless, the basic idea should be translatable to a discrete type space, but the presentation would be quite different and could not be integrated with the current exposition using continuous-time control concepts. From a mathematical point of view, the question of being able to establish BIC or not is one about nonlinear controllability given quite particular control constraints.

Our application of a variable-price auction illustrates the potential benefits of a systematic use of an announcement-contingent mechanism parametrization. It becomes apparent from our analysis that there may be many such truthful mechanism parametrizations, so that selecting the best one from the principal's point of view is subject to further optimization by the principal. Solving problem (P) can be considered an important step on the way to finding optimal mechanisms.

## Acknowledgements

The authors wish to thank Paul Milgrom, an anonymous reviewer, and participants of the 2006 NSF/NBER Decentralization Conference in Paris, France, for helpful comments. Research in part supported by a David Morgenthaler II Faculty Scholar Award.

## References

- Arrow, K.J., 1979. The property rights doctrine and demand revelation under incomplete information. In: Boskin, M. (Ed.), *Economics and Human Welfare*. Academic Press, New York, NY, pp. 23–39.
- Bapna, A., Weber, T.A., 2005. Efficient dynamic allocation with uncertain valuations. Working Paper. Department of Management Science and Engineering, Stanford University.
- Clarke, E.H., 1971. Multipart pricing of public goods. *Public Choice* 2 (1), 19–33.
- d'Aspremont, C., Gérard-Varet, L., 1979. Incentives and incomplete information. *Journal of Public Economics* 11 (1), 25–45.
- Gibbard, A., 1973. Manipulation of voting schemes: a general result. *Econometrica* 41 (4), 587–601.
- Groves, T., 1973. Incentives in teams. *Econometrica* 41 (4), 617–631.
- Holt, C., 1980. Competitive bidding for contracts under alternative auction procedures. *Journal of Political Economy* 88 (3), 433–445.

- Karlin, S., Rinott, Y., 1980. Classes of orderings of measures and related correlation inequalities. I. Multivariate totally positive distributions. *Journal of Multivariate Analysis* 10 (4), 467–498.
- Milgrom, P.R., Weber, R.J., 1982. A theory of auctions and competitive bidding. *Econometrica* 50 (5), 1089–1122.
- Monderer, D., Tennenholtz, M., 2003. k-Price auctions: revenue inequalities, utility equivalence, and competition in auction design. Working Paper. Technion, Haifa, Israel.
- Myerson, R.B., 1979. Incentive compatibility and the bargaining problem. *Econometrica* 47 (1), 61–73.
- Myerson, R.B., 1981. Optimal auction design. *Mathematics of Operations Research* 6 (1), 58–73.
- Palfrey, T.R., Srivastava, S., 1993. Bayesian Implementation. Harwood Academic Publishers, Chur, CH.
- Palfrey, T.R., 2002. Implementation theory. In: Aumann, R., Hart, S. (Eds.), *Handbook of Game Theory*, vol. 3. Elsevier Science, Amsterdam, NL, pp. 2271–2326.
- Postlewaite, A., Schmeidler, D., 1986. Implementation in differential information economies. *Journal of Economic Theory* 39 (1), 14–33.
- Rudin, W., 1976. *Principles of Mathematical Analysis*. McGraw-Hill, New York, NY.
- Vickrey, W., 1961. Counterspeculation, auctions, and competitive sealed tenders. *Journal of Finance* 16 (1), 8–37.
- Weber, T.A., 2005. Screening with externalities. Working Paper. Department of Management Science and Engineering, Stanford University.