
Deterministic Mechanisms, the Revelation Principle, and Ex-Post Constraints

Felix Jarman (University of Mannheim)
Vincent Meisner (Technical University Berlin)

Discussion Paper No. 32

May 20, 2017

Deterministic mechanisms, the revelation principle, and ex-post constraints

Felix Jarman* Vincent Meisner†

May 18, 2017

Abstract

This note establishes a revelation principle in terms of payoff for deterministic mechanisms under ex-post constraints: the maximal payoff implementable by a feasible deterministic mechanism can also be implemented by a feasible deterministic direct mechanism.

JEL-Classification: D82.

Keywords: Mechanism Design, Revelation Principle, Ex-post Constraints.

*University of Mannheim, Department of Economics, Chair for Economic Policy, D-68131 Mannheim, Germany, felix.jarman@gmail.com.

†Technical University Berlin, Department of Economics and Management, Straße des 17. Juni 135, D-10623 Berlin, Germany, vincent.meisner@tu-berlin.de.

1 Introduction

In the analysis of mechanism design problems, economic literature often restricts attention to deterministic mechanisms. In applications, stochastic mechanisms are often deemed problematic as they require that the mechanism designer has access to a credible randomization device which can be implausible or, alternatively, may be prone to manipulation. Therefore, the optimal *deterministic* mechanism is of particular interest. We establish that in ex-post constrained setups a form of the revelation principle also applies to the set of deterministic mechanisms.

As shown by, e.g., Strausz (2003), the classical revelation principle (e.g., Myerson, 1979) does not hold with respect to deterministic mechanisms: there are social choice functions that can be implemented by deterministic *indirect* mechanisms but that cannot be implemented by a deterministic *direct* mechanism in which agents truthfully reveal their type. This failure is due to the possibility that agents play a mixed-strategy equilibrium in the game induced by a deterministic indirect mechanism. In this equilibrium, a non-deterministic social choice function can be implemented, even though the underlying mechanism is deterministic. While a stochastic direct mechanism can replicate players' equilibrium randomization, a deterministic direct mechanism cannot do so by definition.

Strausz (2003) formulates a revelation principle in terms of payoff for the one-agent case under interim constraints. We show that such a form of the revelation principle also holds for the multiple-agents case under ex-post constraints: despite the failure of the revelation principle, it is still without loss of generality to neglect indirect mechanisms if the objective is to identify a social choice function that

- (a) maximizes the expectation of a general payoff function over outcomes (e.g., social welfare or revenue),
- (b) is implementable in dominant strategies, and
- (c) satisfies all additional constraints (if there are any) ex-post.¹

Any deterministic mechanism that maximizes the designer's payoff corresponds to a feasible truthful direct deterministic mechanism that is payoff-equivalent for agents and the designer. Hence, when studying deterministic mechanisms under constraints that have to be satisfied regardless of the strategy of other players, there is no loss of generality when restricting attention to direct truthful mechanisms in optimal mechanism design.

Depending on the application, imposing constraints ex-post instead of interim or ex-ante is more suitable. We use our result in Jarman and Meisner (2016), where a procurer can only spend a *fixed* budget, i.e., the budget must suffice for any outcome of the mechanism. Alternatively, limited liability or the possibility of default motivate an ex-post participation constraint. Dominant-strategy implementable mechanisms are popular, as they are easy to explain and are not prone to manipulation or the misspecification of beliefs. When considering such mechanisms for the latter reason, it is indeed consistent not to consider stochastic mechanisms because they require agents (who may believe in other priors) to believe in the mechanism designers' randomization distribution.

¹There can be additional constraints that are imposed ex-ante or interim, if these constraints are implied by ex-post constraints.

The appeal of both deterministic and dominant-strategy incentive-compatible (DIC) mechanisms has led to many contributions on mechanism equivalence. For any Bayesian incentive-compatible (BIC) mechanism, Chen, Wei, Li, and Sun (2016) construct an equivalent *deterministic* BIC mechanism delivering the same interim outcome.² Gershkov, Goeree, Kushnir, Moldovanu, and Shi (2013) show that for any feasible BIC mechanism there exists a DIC mechanism that yields the same interim expected payoffs. Both equivalence results do not generalize our result because the equivalent mechanisms can feature other ex-post outcomes that may violate our ex-post constraints. Moreover, our constructed direct mechanism is not equivalent to, but payoff-dominates, a corresponding indirect mechanism.

2 Model

A (mechanism) designer faces a set of N agents. A type profile $\theta = (\theta_1, \dots, \theta_N)$ is drawn from a finite type space $\Theta = \Theta_1 \times \dots \times \Theta_N$ according to some probability mass function (pmf) P , and each agent i is privately informed about his type θ_i .

A social choice function f maps type profiles into distributions over outcomes,

$$f : \Theta \rightarrow \Delta X,$$

where X is a finite set of outcomes. Let $f_\theta : X \rightarrow [0, 1]$ denote the pmf representing the randomization over outcomes corresponding to $f(\theta)$. If f_θ is degenerate for all $\theta \in \Theta$, we call the corresponding social choice function deterministic. Let \mathcal{F}_d and \mathcal{F}_s denote the set of deterministic and stochastic social choice functions, respectively, where $\mathcal{F}_d \subset \mathcal{F}_s$. The designer wants to elicit information θ to select a distribution over outcomes maximizing her expected payoff,

$$\mathbb{E}_\theta [\omega(\theta, f)] = \mathbb{E}_\theta [\mathbb{E}_x [w(\theta, x) | \theta, f]] = \sum_{\theta \in \Theta} \sum_{x \in X} w(\theta, x) f_\theta(x) P(\theta),$$

where $w(\theta, x)$ denotes her ex-post payoff when type profile θ and outcome x realize. The following three paragraphs elaborate on her restrictions in doing so.

(I) She can only propose deterministic mechanisms $M = (A, g)$ consisting of a collection of action spaces $A = A_1 \times \dots \times A_N$ and a *deterministic* outcome function $g : A \rightarrow X$ that maps action profiles $a = (a_1, \dots, a_N) \in A$ into outcomes. To ensure the existence of an equilibrium in the game induced by the mechanism, we only consider mechanisms with finite action spaces. The set of direct mechanisms consists of all mechanisms such that $A = \Theta$, i.e., all mechanisms that simply ask the agents to report their type.

(II) A type- θ_i agent's ex-post utility from outcome x is denoted by $U_i(x, \theta_i)$. Let a type- θ_i agent i 's payoff from playing action $a_i \in A_i$ against actions $a_{-i} \in A_{-i} = \times_{j \neq i} A_j$ in mechanism M be denoted by $u_i^M(a_i, a_{-i}, \theta_i) = U_i(g(a_i, a_{-i}), \theta_i)$. While a strategy maps types into distributions over actions, we denote it by a pmf $\sigma_i : \Theta_i \times A_i \rightarrow [0, 1]$ and let

²Essentially, their novel methodology of “mutual purification” replicates a stochastic outcome for one agent by exploiting the inherent randomness of another agent's atomless type draw. Despite a restriction to atomless type distributions, their setting is very general.

Σ_i be the space of all such pmfs. Let $\sigma = (\sigma_1, \dots, \sigma_N) = (\sigma_i, \sigma_{-i})$. A strategy σ_i is a (weakly) dominant strategy in the game induced by mechanism M if for all $\theta_i \in \Theta_i$,

$$\sum_{b_i \in A_i} u_i^M(b_i, a_{-i}, \theta_i) \sigma_i(\theta_i, b_i) \geq \sum_{b_i \in A_i} u_i^M(b_i, a_{-i}, \theta_i) \sigma'_i(\theta_i, b_i) \quad \forall \sigma'_i \in \Sigma_i, \forall a_{-i} \in A_{-i}, \quad (1)$$

and the inequality holds strictly for at least one opponent action profile a_{-i} .³ A mechanism M implements social choice function f in dominant strategies if

$$f_\theta(x) = \sum_{a \in A} \mathbb{I}_{\{g(a)=x\}} \sigma(\theta, a) \quad \forall \theta \in \Theta, \quad (2)$$

where σ constitutes an equilibrium in (weakly) dominant strategies in the game induced by M , i.e., (1) holds for all $\sigma_i \in \sigma$.

(III) Depending on the application, the designer can be restricted further by additional feasibility constraints, such as individual rationality or budget balance. We say that the designer faces ex-post constraints if these constraints must be satisfied regardless of the strategies other agents play. The resulting set of social choice functions that are ex-post constrained implementable by a (direct or indirect) deterministic mechanism is given by $\mathcal{F}^* \subset \mathcal{F}_s$.

Definition 1. A (potentially stochastic) social choice function f is ex-post DDS-implementable, $f \in \mathcal{F}^*$, if (I) there exists a Deterministic mechanism M that (II) implements f in Dominant Strategies, while (III) all other (if there are any) constraints are satisfied ex-post.

Consequently, the designer's optimization problem is

$$\max_f \mathbb{E}_\theta[\omega(\theta, f)] \quad \text{such that } f \in \mathcal{F}^*. \quad (3)$$

3 A revelation principle in terms of payoff

Proposition 1. For any stochastic social choice function $f \in \mathcal{F}_s$ that is ex-post DDS-implementable by an indirect mechanism, there exists a deterministic social choice function $\hat{f} \in \mathcal{F}_d$ that

1. is ex-post DDS-implementable, $\hat{f} \in \mathcal{F}^*$, by a direct mechanism, and
2. weakly dominates f state-by-state in terms of the designer's payoff,

$$\omega(\theta, \hat{f}) \geq \omega(\theta, f) \quad \forall \theta \in \Theta.$$

Proof. Suppose $f \in \mathcal{F}_s$ and $f \notin \mathcal{F}_d$. Let the deterministic indirect mechanism that DDS-implements f be denoted by $M = (A, g)$. Given an arbitrary type profile θ , consider a strategy profile mixing over (pure-strategy) action profiles $\hat{a}_\theta = (\hat{a}_{\theta_1}, \dots, \hat{a}_{\theta_N}) \in \hat{A}_\theta \subseteq A$.

³Since σ_i is weakly better against all pure strategies a_{-i} , it is also weakly better against any convex combination thereof, i.e., against all mixed strategy profiles.

Because each agent i plays a dominant strategy, every pure strategy \widehat{a}_{θ_i} over which he randomizes must be a weakly dominant strategy, too, and thus

$$\sum_{b_i \in A_i} u_i^M(b_i, a_{-i}, \theta_i) \sigma_i(\theta_i, b_i) = u_i^M(\widehat{a}_{\theta_i}, a_{-i}, \theta_i) \quad \forall \widehat{a}_{\theta_i} \in \widehat{A}_{\theta_i}, \forall a_{-i} \in A_{-i}.$$

Otherwise (1) would be violated, as the strategy must be a best response for agent i regardless of the other agents' strategies.

The designer's payoff given type profile θ can be stated as

$$\begin{aligned} \omega(\theta, f) &= \sum_{x \in X} w(\theta, x) f_\theta(x) \\ &\stackrel{\text{by (2)}}{=} \sum_{x \in X} w(\theta, x) \sum_{a \in A} \mathbb{I}_{\{g(a)=x\}} \sigma(\theta, a) \\ &= \sum_{x \in X} \sum_{a \in \widehat{A}_\theta} w(\theta, x) \mathbb{I}_{\{g(a)=x\}} \sigma(\theta, a) \\ &= \sum_{\widehat{a}_\theta \in \widehat{A}_\theta} w(\theta, g(\widehat{a}_\theta)) \sigma(\theta, \widehat{a}_\theta). \end{aligned}$$

Define action profile $\bar{a}_\theta = (\bar{a}_{\theta_1}, \dots, \bar{a}_{\theta_N})$ such that

$$\bar{a}_\theta \in \arg \max_{a \in \widehat{A}_\theta} w(\theta, g(a)).$$

That is, \bar{a}_θ represents the designer's most preferred pure-strategy profile over which the mixed strategy randomizes. Following the same argument as above, \bar{a}_θ is an equilibrium strategy profile in weakly dominant pure strategies for type profile θ in mechanism M . Similarly, any outcome that can result from f for any type profile θ must be ex-post feasible. Therefore, $g(\bar{a}_\theta)$ is feasible as well.

Define \widehat{f} such that $\widehat{f}_\theta(x) = \mathbb{I}_{g(\bar{a}_\theta)=x}$ for any type profile θ for which σ is not degenerate. By construction, \widehat{f} generates a weakly higher payoff for any type profile,

$$\omega(\theta, \widehat{f}) = w(\theta, g(\bar{a}_\theta)) \geq \omega(\theta, f) \quad \forall \theta \in \Theta.$$

Because \widehat{f} is ex-post feasible and deterministic, $\widehat{f} \in \mathcal{F}^* \cap \mathcal{F}_d$, it is DDS-implementable by a direct mechanism. \square

Proposition 1 states that for any stochastic ex-post DDS-implementable social choice function there exists a DDS-implementable deterministic social choice function which satisfies the same ex-post constraints and gives the designer a weakly larger payoff. The result can be interpreted as a variation of the revelation principle formulated in terms of payoff: While not every social choice function that is DDS-implementable by an indirect mechanism is also DDS-implementable by a direct mechanism, there exists a direct deterministic mechanism that yields a weakly larger payoff for the designer. Therefore, the optimal DDS-implementable social choice function can always be DDS-implemented by a direct mechanism when all other constraints (if there are any) must hold ex-post. Put differently, without loss of generality, we can replace (3) with

$$\max_f \mathbb{E}_\theta[w(\theta, f)] \quad \text{such that } f \in \mathcal{F}^* \cap \mathcal{F}_d. \quad (4)$$

Strausz (2003) obtains a similar result for private-value environments with a single agent under interim constraints. With a single agent, the agent’s best response trivially is a dominant strategy. Similarly, if the agent’s participation constraint in a deterministic mechanism holds interim, it also holds ex-post. Strausz (2003) provides an example with two agents such that his revelation principle in terms of payoff fails: he imposes an interim participation constraint (individual rationality), and mixing in the indirect mechanism guarantees the agents their reservation utility.

In contrast, a participation constraint in our setting would have to hold ex post, i.e., agents must obtain at least their reservation utility regardless of the other agents’ strategies. Therefore, agents cannot play a mixed strategy that attaches positive weight to a pure strategy that could, against any possible strategies of the other agents, yield a payoff that is less than the reservation utility.

4 Conclusion

It is known that the classical revelation principle fails when attention is restricted to deterministic mechanisms. We establish that a mechanism designer can restrict attention to direct mechanisms when all constraints have to be satisfied, regardless of which strategies the other agents play.

5 Acknowledgments

We thank Thomas Schacherer, Roland Strausz, and an anonymous referee for feedback. Support by Deutsche Forschungsgesellschaft through CRC 884 (Felix Jarman) and CRC TRR 190 (Vincent Meisner) is gratefully acknowledged.

References

- CHEN, Y.-C., H. WEI, J. LI, AND Y. SUN (2016): “Equivalence of Stochastic and Deterministic Mechanisms,” *Mimeo, National University of Singapore*.
- GERSHKOV, A., J. K. GOEREE, A. KUSHNIR, B. MOLDOVANU, AND X. SHI (2013): “On the equivalence of Bayesian and dominant strategy implementation,” *Econometrica*, 81(1), 197–220.
- JARMAN, F., AND V. MEISNER (2016): “Ex-post Optimal Knapsack Procurement,” *Mimeo, University of Mannheim*.
- MYERSON, R. B. (1979): “Incentive Compatibility and the Bargaining Problem,” *Econometrica*, 47(1), 61–74.
- STRAUSZ, R. (2003): “Deterministic mechanisms and the revelation principle,” *Economics Letters*, 79(3), 333–337.