Bayesian-Nash games*

Sergei Izmalkov
14.129 Advanced Contract Theory

February 23, 2005

1 A general incomplete information setting

1.1 Primitives

A finite group of *players* (economic agents), denoted by $N = \{1, \ldots, n\}$, interact. Any interaction can be represented as a simultaneous non-cooperative choice of individual plans of action.[1] A set of possible decisions or choices (a set of all possible contingent plans of action) is a set of *strategies*. Denote S_i to be a set of strategies available to player i ($s_i \in S_i$ is a generic element); $s = (s_1, \ldots, s_n) = \{s_i\}_{i \in N}$ is a profile of strategies of all the players, $s \in S = \times_{i \in N} S_i$; $s_{-i} = s \smallsetminus \{s_i\}$, $S_{-i} = S \setminus S_i$, $s = (s_i, s_{-i})$. A set of *alternatives* (allocations), A, is a set of all possible outcomes. A *mechanism* is a rule that for any collection of strategies selects a probability distribution over the set of alternatives A, $M : S \to \Delta(A)$.

Example 1: *Voting.* A set of alternatives A is a set of possible candidates to be chosen from. Players are individuals with a right to vote. Sets of strategies are determined by a voting procedure. In a simple ballot voting, for example, each player submits a ballot for some candidate, so $S_i = A$. In a multistage voting, like a procedure to select a city to hold Olympics, a strategy has to name a city in the first round of voting, a city in the second round of voting conditional on results of the first round, and so on, for each round conditional on previous results. A voting mechanism specifies how the winner is selected. For instance, in a simple majority voting, bar ties, $M(s) = \arg\max_{a \in A} \#_{i \in N}\{s_i = a\}$.

Example 2: *Auctions.* A set of alternatives A is a set of all possible allocations of goods for sale and transfers involved. For example, with one object for sale and only the winner paying, a set of alternatives is $A = N \times \mathbb{R}_+$, a pair $(w, m) \in A$ says that player $w \in N$ is the winner and has to pay m. A set of strategies depends on an auction format. In sealed-bid auctions, a strategy is a bid, $S_i = \mathbb{R}_+$. In open or dynamic auctions, like an English (a usual ascending price) auction, a strategy has to specify how to act (bid) for any possible scenario that can occur in the auction. The winner is determined according to the rules of the specific format. In the first-price sealed-bid auction the winner is the highest bidder and has to pay own bid. Thus, bar ties, $w = \arg\max_{i \in N} s_i$ and $m = s_w$. In the second-price sealed-bid auction the winner is also the highest bidder but pays the highest bid among the rest of the participants, $m = \max_{j \neq w} s_j$.

Example 3: *Public good provision* (discrete setting). There is a public project that can be built if sufficient funds are collected from citizens (to cover a cost $c \in \mathbb{R}_+$). Thus, an allocation $(b, m_1, \ldots m_n) \in$

*To appear as part of "Secure direct implementation" by Sergei Izmalkov, Matt Lepinsky, and Silvio Micali.

[1] This is called a normal or strategic form representation. For a detailed coverage of game theory and of incomplete information games in particular the reader is referred to two excellent textbooks on the subject, Fudenberg and Tirole (1991) and Osborne and Rubinstein (1997).

$\{0,1\} \times \mathbb{R}^n_+$ specifies whether a project is built ($b = 1$) or not ($b = 0$), and for each i a transfer m_i. The strategies depend on the procedure that is used to decide on whether to build and on contributions. For example, for the case of private voluntary contributions, each player contributes $c_i \in S_i = \mathbb{R}_+$. The corresponding mechanism sets $m_i = c_i$, the project is built, $b = 1$, only if $\sum_{i\in N} m_i \geq c$.

To complete description of a game players' preferences over alternatives have to be specified. In a complete information setting players preferences are commonly known. In an incomplete information setting some players are not certain about preferences of the others. In reality, essentially in any interaction something is not known to all participants. How much money or other resources a player has? What are the costs of production? What does a player know (or think) about what others know? Often, the answers to these questions are only known to the player in focus.

Following Harsanyi (1967-68), players' uncertainty is added as follows. A special player, *Nature*, moves first and selects a profile of types, $t = \{t_i\}_{i\in N}$, $t_i \in T_i$, for each of the players according to some commonly known distribution p over $T = \times_{i\in N} T_i$. A *type* is a complete description of all relevant characteristics of a given player. A player observes her type, beliefs about the types of the others, $t_{-i} = t \setminus \{t_i\}$, are calculated using conditional distribution $p(t_{-i}|t_i)$.[2] To complete specification, each players' preferences over lotteries over $T \times A$ need to be defined. We will assume that these preferences satisfy expected utility axioms of von Neumann and Morgenstern, thus it will suffice to specify payoff functions, $u_i : T \times A \to \mathbb{R}$ for each i. Altogether, a game of incomplete information, or *Bayesian* game, is described by a septuple, $\Gamma = (N, \{T_i\}_{i\in N}, \{S_i\}_{i\in N}, A, M, \{u_i\}_{i\in N}, p)$.[3]

Note that essentially an incomplete information game can be thought of as a very large game with complete but imperfect information. Players in that game are all possible player-type combinations, and, given any selected profile of strategies, a payoff to player (i, t_i) is $\sum_{t_{-i}} p(t_{-i}|t_i)u_i(t, M(s))$.

Definition: A *pure strategy* of player i is a function $s_i(t_i)$ that for each $t_i \in T_i$ selects an element of S_i. A *mixed strategy* of player i is a function $\sigma_i(t_i)$ that for each $t_i \in T_i$ selects a probability distribution over S_i. We denote $s(t) = (s_1(t_1), \ldots s_n(t_n))$, $s_{-i}(t_{-i}) = s(t) \setminus \{s_i(t_i)\}$, mixed profiles $\sigma(t)$ and $\sigma_{-i}(t_{-i})$ are defined similarly.

Without any confusion we can define $u_i(t, s) = u_i(t, M(s))$, a payoff to a mixed $\sigma(t)$ is

$$u_i(t, \sigma(t)) = E_{\sigma_1(t_1),\ldots,\sigma_n(t_n)} u_i(t_1, \ldots t_n, s_1, \ldots, s_n).$$

Definition: A selection of (mixed) strategies $\{\sigma_i^*(t_i)\}_{i\in N}$ is:
a *Bayesian-Nash equilibrium* of game Γ if, for each type t_i of player i, for any $\sigma_i(t_i)$,

$$E_{t_{-i}} u(t_i, \sigma_i^*(t_i), t_{-i}, \sigma_{-i}^*(t_{-i})) \geq E_{t_{-i}} u(t_i, \sigma_i(t_i), t_{-i}, \sigma_{-i}^*(t_{-i})); \tag{1}$$

a *dominant strategy equilibrium* of game Γ if, for each type t_i of player i, for any $\sigma_i(t_i)$ and any $s_{-i}(t_{-i})$,[4]

$$u(t_i, \sigma_i^*(t_i), t_{-i}, s_{-i}(t_{-i})) \geq u(t_i, \sigma_i(t_i), t_{-i}, s_{-i}(t_{-i})); \tag{2}$$

[2] This is sometimes referred as a reduced form construction. An alternative description starts with states of nature, $\theta \in \Theta$. A type is interpeted as a privately observed signal about θ. Types are generated according to signal functions $\tau_i : \Theta \to T_i$. For a more detailed treatment see Osborne and Rubinstein (1997).

[3] To define a Bayesian game it suffices to specify N, T, S, p, and payofs over lotteries over $T \times S$ (which is a typical approach). A restriction to mechanisms is without loss of generality, since we can always set $A = S$, and $M(s) = s$. Alternatives and mechanisms are explicitly introduced for two purposes: to have a more convenient language and to describe explicitly what is publicly observable (a resulting allocation, e.g.) and what remains private (actual strategies, e.g.).

[4] It suffices to check dominant strategy incentive constraints only against pure strategies of opponents. Existence of mixed strategies that violate the constraints implies existence of pure strategies that do the same. Also note that any pure strategy s_i in the support of $\sigma_i^*(t_i)$ is a (weakly) dominant strategy for player i of type t_i. Thus, any selection of pure strategies from the support of mixed dominant strategy equilibrium will form a pure dominant strategy equilibrium.

an *ex post equilibrium* of game Γ if, for each type t_i of player i, for any $\sigma_i(t_i)$, for all t_{-i},

$$u(t_i, \sigma_i^*(t_i), t_{-i}, \sigma_{-i}^*(t_{-i})) \geq u(t_i, \sigma_i(t_i), t_{-i}, \sigma_{-i}(t_{-i})). \tag{3}$$

In words, Bayesian-Nash incentive constraints (1) require a best response against best responses of all the players of all possible types, and are the weakest type of constraints. Dominant strategy incentive constraints (2) are the strongest and require that a given strategy is a best response against any selection of the strategies for any realization of the opponents' types. Ex post incentive constraints (3) fall in between, selected strategies must remain Nash equilibrium strategies for each realization of types. Alternatively, equilibrium conditions must hold *ex post*, even if realized types for all the players become commonly known. Both dominant strategy and ex post equilibria are distribution-free, they do not depend on a given specification of distribution of types, p. Lastly, any dominant strategy equilibrium is an ex post equilibrium, any ex post equilibrium is Bayesian-Nash equilibrium.

Example 2 (revisited). Consider an auction setting from above (one object, N bidders). The most common paradigm in which it is studied is a symmetric independent private values paradigm.[5] Each bidder learns her own value V_i (this is her type) — a maximal willingness to pay for the object. Valuations of all the bidders are identically and independently distributed, and only player i knows her V_i. The simplest case is $V_i \sim U[0,1]$; thus, $T = [0,1]^n$. For an allocation (w, m) payoffs are defined as follows: $u_w(w, m) = V_w - m$, and for any $j \neq w$, $u_j = 0$.

The first-price sealed bid auction admits a unique Bayesian-Nash equilibrium. In it, each player i follows the same strategy, $b^I(V_i) = \frac{n-1}{n} V_i$.[6] This equilibrium is neither in dominant strategies nor ex post. The second-price sealed bid auction has multiple Bayesian-Nash equilibria (and ex post equilibria). It does admit a dominant strategy equilibrium. In it, each player simply bids her own value, $b^{II}(V_i) = V_i$. Because of the strong incentive constraints and weak informational requirements (players need not to know anything about preferences of the others) this equilibrium is usually selected as a solution to the second-price auction.

1.2 Direct mechanisms and revelation principle

Definition: A game is called a *direct game* if the set of strategies for each player coincides with the set of possible types, $S_i = T_i$. A *direct mechanism*, is a mechanism of a direct game.

Theorem 1 (Revelation Principle) [7] *For any Bayesian-Nash (dominant strategy, ex post) equilibrium $\{\sigma_i^*(t_i)\}_{i\in N}$ of a Bayesian game $\Gamma = (N, \{T_i\}_{i\in N}, \{S_i\}_{i\in N}, A, M, \{u_i\}_{i\in N}, p)$ there exists a direct game $\Gamma^D = (N, \{T_i\}_{i\in N}, \{T_i\}_{i\in N}, A, M^D, \{u_i\}_{i\in N}, p)$ such that* truthtelling, *$\{s_i^*(t_i) = t_i\}_{i\in N}$,*

1. *is a Bayesian-Nash (dominant strategy, ex post) equilibrium;*
2. *is* outcome equivalent *to the equilibrium of the original (indirect) game: for any $t \in T$, $M(\sigma^*(t)) = M^D(t)$ (as probability distributions over A).*

Proof. Define $M^D(t) = M(\sigma^*(t))$. Validity of incentive constraints (1) (or (2) for dominant strategies, (3) for ex post constraints) for $s^* = t$ of Γ^D trivially follows from the corresponding constraints for Γ. ■

The brilliantly simple logic of the revelation principle can be summarized as follows. Imagine that instead of playing the game in person each player can ask a trusted agent to play it in her place. Imagine that a

[5] For a comprehensive coverage of auction theory, including facts presented here, the reader is referred to Krishna (2002).

[6] An equilibrium bid b of player i with value V solves $\max_b (V - b)\Pr(b > \max_{j\neq i} b^I(V_j))$.

[7] Early formulations and utilizations of reveleation priciple appear in Gibbard (1973), Green and Laffont (1977), Dasgupta, Hammond, and Maskin (1979), Myerson (1979), Myerson (1981).

single *trusted party* is this trusted agent for each type of each player. Having a particular equilibrium in mind, a player needs only to tell the trusted party her type. If a trusted party for each type of player i follows an equilibrium strategy of that type, $\sigma_i^*(t_i)$, then it is optimal to make a truthful announcement. Indeed, if player i of type t_i tells the truth, she trusts the trusted party to select and follow $\sigma_i^*(t_i)$ on her behalf. If she reports a different type, perhaps a different strategy would be selected. Clearly, if $\sigma_i^*(t_i)$ is the best strategy for type t_i then a different strategy would not be better, so telling the truth is the best strategy.

The most notable application of revelation principle—a cornerstone of mechanism design and implementation theory—is to answer questions: whether an allocation rule satisfying certain properties can be implemented in principle. In the context of Bayesian games, an *allocation rule*, or a *social choice rule*, is a function that given a profile of types selects an alternative. To *implement* the allocation rule is to find a game (strategy sets and a mechanism) that has an equilibrium that for each possible profile of types results in the alternative prescribed by the allocation rule (sometimes it is required that all equilibria must implement the allocation rule). The strength of equilibrium incentive constraints determines the strength of implementation (Bayesian-Nash, ex post, or in dominant strategies).

By revelation principle, if an allocation rule is implemented by some game, then it is also implemented truthfully by the corresponding direct game, and so the (direct) mechanism of the direct game coincides with the allocation rule. Thus, to verify whether a given allocation can be implemented it suffices to check whether it satisfies equilibrium incentive constraints. If it can be shown that incentive constraints are necessarily violated, a given allocation rule cannot be implemented. If they are satisfied, a direct game with the mechanism equal to the allocation rule is a game that implements it.

References

Dasgupta, P. S., P. J. Hammond, and E. S. Maskin (1979): "The Implementation of Social Choice Rules: Some General Results on Incentive Compatibility," *Review of Economic Studies*, 46(2), 185–216.

Fudenberg, D., and J. Tirole (1991): *Game Theory*. MIT Press, Cambridge, Massachusetts.

Gibbard, A. (1973): "Manipulation of Voting Schemes: A General Result," *Econometrica*, 41(4), 587–601.

Green, J., and J.-J. Laffont (1977): "Characterization of Satisfactory Mechanisms for the Revelation of Preferences for Public Goods," *Econometrica*, 45(2), 427–38.

Harsanyi, J. (1967-68): "Games with Incomplete Information Played by Bayesian Players," *Management Science*, 14, 159–182,320–334,486–502.

Krishna, V. (2002): *Auction Theory*. Academic Press, San Diego, California.

Myerson, R. (1981): "Optimal Auction Design," *Mathematics of Operations Research*, 6, 58–73.

Myerson, R. B. (1979): "Incentive Compatibility and the Bargaining Problem," *Econometrica*, 47(1), 61–73.

Osborne, M. J., and A. Rubinstein (1997): *Game Theory*. MIT Press, Cambridge, Massachusetts.