Lecture Notes for Game Theory Tutorials

at

Institute for Mathematical Studies

National University of Singapore

by

Sudhir A. Shah*

Department of Economics

Delhi School of Economics

June 2005

* These notes are preliminary and incomplete. I will be grateful for feedback regarding errors and possible improvements. All correspondence should be addressed to: Department of Economics, Delhi School of Economics, University of Delhi, Delhi 110 007, India. Telephone: (+91)(11) 2766-7005. Fax: (+91)(11) 2766-7159. E-mail: sudhir@econdse.org

Contents

1. Representations and equilibria of games

1.1 Extensive form representation

In this section we set up the extensive form model used to represent and study games.[1] This model explicitly displays the following data:

(a) the set of players,

(b) the order of moves in the game,

(c) for every decision point, the identity of the player who is to move,

(d) the information a player has when asked to move,

(e) the menu of actions available to a player when asked to move,

(f) the payoffs for all players when the game terminates, and

(g) the probabilities associated with random moves.

More formally, we have

Definition 1.1.1. *A collection* $\Gamma(e) = \{(T, \succ), N, \iota, (D_t)_{t\in X}, S, v, \rho\}$ *is called a game in extensive form.*

$(T, \succ)$ is called the *game tree.* T is a set whose elements are called *nodes.* $\succ$ is a binary relation on T called the *precedence relation.* $\succ$ orders the nodes and is used to describe the order of moves in the game. For nodes t and t', $t \succ t'$ is interpreted as t' *precedes* t, or equivalently, t *follows* t'. The set of *predecessors* of node t is denoted by $\succ^t = \{t' \in T \mid t \succ t'\}$ and the set of *successors* of t is denoted by $\succ_t = \{t' \in T \mid t' \succ t\}$. The set of *initial nodes* is $W = \{t \in T \mid \succ^t = \emptyset\}$ and the set of *terminal nodes* is $Z = \{t \in T \mid \succ_t = \emptyset\}$. The set of *decision nodes* is $X = T - Z$. The set of terminal nodes that can be reached from node t is

$$Z_t = \begin{cases} Z \cap \succ_t, & \text{if } t \in X \\ \{t\}, & \text{if } t \in Z \end{cases}$$

The set of *maximal elements* of $T' \subset T$ is $\max T' = \{t \in T' \mid \succ_t \cap T' = \emptyset\}$ and the set of *minimal elements* of $T' \subset T$ is $\min T' = \{t \in T' \mid \succ^t \cap T' = \emptyset\}$. The set of *immediate predecessors* of node t is $b(t) = \max \succ^t$ and the set of *immediate successors* of t is $a(t) = \min \succ_t$. A sequence of nodes $(t_1, \ldots, t_m)$ is said to be the *path* from t_1 to t_m

[1] The first systematic account of game theory is contained in von Neumann and Morgenstern (1944). The extensive form model described below was formulated in Kuhn (1953).

if $t_i = b(t_{i+1})$ for $i = 1, \ldots, m-1$. A path $(t_1, \ldots, t_m)$ is said to be a *play* of the game if $t_1 \in W$ and $t_m \in Z$.

Assumption 1.1.2. *Given* $(T, \succ)$,

(a) T *is a finite set.*

(b) $(T, \succ)$ *is an arborescence, i.e.,* $\succ$ *is transitive, asymmetric, and for all* t, t', $t'' \in T$,

$$\langle t'' \succ t \quad \wedge \quad t'' \succ t' \quad \wedge \quad t \neq t' \rangle \qquad \Longrightarrow \qquad \langle t \succ t' \quad \vee \quad t' \succ t \rangle$$

(c) $W = \{t_0\}$.

It is easy to check that Assumption 1.1.2 implies that $b(t)$ is a singleton for every $t \in T - \{t_0\}$. Also, for every $t \in T$, there is a unique path from t_0 to t.

N is the set of players. If $0 \in N$, then 0 is referred to as *Nature*.

Assumption 1.1.3. $1 \leq |N| < \infty$.

$\iota : X \to N$ is the assignment of decision nodes to the various players. When $\iota(t) = j$, we say that *node* t *belongs to player* j. The set $\iota^{-1}(\{i\})$ is called player i's *player set*. The family of player sets, $\{\iota^{-1}(\{i\}) \mid i \in N\}$, form a partition of X, called the *player partition*. D_t is the set of decisions (or actions) available at decision node $t \in X$.

Assumption 1.1.4. *For every* $t \in X$, *there exists a bijection* $\alpha_t : D_t \to a(t)$.

Thus, at every decision node t, every decision is associated with one and only one immediate successor node. As T is finite, $|D_t| < \infty$ for every $t \in T$. We now come to the specification of information in a game. Let S be a partition of X. S is called the *information partition*, and the generic element of S, denoted by s, is called an *information set*. The information set containing the node t is denoted by $s(t)$.

Assumption 1.1.5. *For every* $s \in S$ *and all* t, $t' \in s$

(a) $t \notin \succ^{t'}$ *and* $t' \notin \succ^{t}$,

(b) $\iota(t) = \iota(t')$, *and*

(c) $D_t = D_{t'}$.

Condition (b) implies that the information partition is a refinement of the player partition; for every $s \in S$, there exists $i \in N$ such that $s \subset \iota^{-1}(\{i\})$. $S_i = \{s \in S \mid s \subset \iota^{-1}(\{i\})\}$ is the collection of *information sets belonging to player* i; we shall often abuse notation by using $\iota(s) = i$ to denote the fact that $s \in S_i$.[2] The interpretation of an information set $s \in S_i$ is that the nodes in s are indistinguishable, given player i's information when he is required to move from one of the nodes in s. Given this interpretation, condition (a) is quite natural. It requires that if two nodes are indistinguishable, then neither can be a predecessor of the other: players remember the nodes where they have already made decisions. Condition (c) requires that the actions available at all nodes in a given information set be the same; otherwise, the player would be able to distinguish nodes by the actions available at the nodes. Assumption 1.1.5(c) allows us to set $D_s = D_t$, where $t \in s$. For a given information set $s \in S$, the set of successor nodes is denoted $\succ_s = \cup_{t \in s} \succ_t$ and the set of predecessor nodes is denoted $\succ^s = \cup_{t \in s} \succ^t$. The set of terminal nodes reachable from an information set s is $Z_s = \succ_s \cap Z$. $v : N - \{0\} \times Z \to \Re$ is the assignment of payoffs at terminal nodes to the various players, other than Nature. $v_i : Z \to \Re$ is the utility function of player $i \in N - \{0\}$ defined by the formula $v_i(t) = v(i, t)$, $t \in Z$. We interpret v_i to be player i's von Neumann-Morgenstern utility function.[3]

It often is necessary to introduce *exogenous* randomness in various games. We model this randomness as resulting from Nature's moves. The difference between Nature and the other players is that Nature's actions are non-purposive and are given exogenously as part of the description of the game.

Assumption 1.1.6. $\iota(t) = 0$ *implies* $t = t_0$.

If there is a Nature move in the game, it is at the root node. If $\iota(t_0) = 0$, then $\rho \in \Delta(D_{t_0})$ specifies the probability distribution of Nature's action. This specification is "without loss of generality" because any finite game in which nature moves more than once can be formally transformed into one in which nature moves only once at the beginning of the game.

Exercise. *Define a notion of equivalence for which the transformation is "without loss of generality"? Outline the procedure for making this transformation.*

---

[2] We are merely identifying the elements of an equivalence class.

[3] See von Neumann and Morgenstern (1944) for the first formulation. A very general existence theorem for such functions is contained in Herstein and Milnor ().

Definition 1.1.7. $\Gamma(e)$ *is a standard game if it satisfies Assumptions 1.1.2 to 1.1.6.*

Definition 1.1.8. *Given a standard game* $\Gamma(e)$ *and* $t^* \in T$*, suppose* $T^* = \{t^*\} \cup \succ_{t^*}$*,* $\succ^*$ *is the restriction of* $\succ$ *to* T^**,* ι^* *is the restriction of* ι *to* $X \cap T^*$*,* $N^* = \iota^*(X \cap T^*)$*,* $S^* = \{s \in S \mid s \cap T^* \neq \emptyset\}$ *and* v^* *is the restriction of* v *to* $N^* \times Z_{t^*}$*. The data,*

$$\Gamma(e, t^*) = \{(T^*, \succ^*), N^*, \iota^*, (D_t)_{t \in T^* \cap X}, S^*, v^*, \rho\}$$

is called a proper subgame of $\Gamma(e)$ *with root node* t^* *if* $s \subset T^*$ *for every* $s \in S^*$*.*

The critical feature of a proper subgame is that it is isolated from the rest of the game by the requirement that no information set should overlap a proper subgame and its complement. Once the play of a game enters a proper subgame at t^*, all players know that the subsequent play will be restricted to $\succ_{t^*}$. Note that $\Gamma(e)$ is the proper subgame $\Gamma(e, t_0)$ of itself.

1.2 Strategic form representation: What is to be Done?

In this section we shall describe various constructions using a standard game $\Gamma(e)$ as data. $C_i = \prod_{s \in S_i} D_s$ is called the set of *player i's (pure) strategies.* A pure strategy is a complete plan of action for the player, which specifies an action, from the available set of actions, at every information set belonging to the player. As $\Gamma(e)$ is finite, C_i is a finite set. As N is finite, the set of *pure strategy profiles* $C = \prod_{i \in N} C_i$ is finite.

A strategy profile $c \in C$ *generates a unique play and a unique terminal node of* $\Gamma(e)$*.* The root node t_0 is controlled by player $\iota(t_0)$. As, by our assumptions, $\{t_0\}$ is an information set, the chosen action is $c_{\iota(t_0)}(\{t_0\})$. Consider $t_n \in X$; t_n belongs to the information set $s(t_n) \in S_{\iota(t_n)}$. As the chosen action at t_n is $c_{\iota(t_n)} \circ s(t_n)$, the immediate successor of t_n is $t_{n+1} = \alpha_{t_n} \circ c_{\iota(t_n)} \circ s(t_n)$. Using this inductive definition, we construct a path $\{t_0, t_1, \ldots, t_k\}$. As T is finite, $t_k \in Z$ for k sufficiently large, meaning that $\{t_0, \ldots, t_k\}$ is the unique play generated by c. Thus, the implied mapping from C to the set of plays is a function. It is easy to see that this function is surjective, i.e., every play is generated by some profile of pure strategies, but not necessarily injective, i.e., there exist games in which distinct strategy profiles generate the same play. Let $\{t_0, \ldots, t_k\}$ be a play of $\Gamma(e)$. It is easy to show that the function $\{t_0, \ldots, t_k\} \mapsto t_k$, which maps plays of $\Gamma(e)$ into Z, is bijective. Consequently, the composition of the above two mappings, from C to Z, is a function; it follows from the properties of the two functions that the composition function is surjective, but not necessarily injective. Let $z(c) \in Z$ be the unique terminal node reached when the strategy profile $c \in C$ is implemented.

$\Delta(C_i)$, the set of probability measures over C_i, is the set of *mixed strategies* for player $i \in N$. As C_i is finite, $\Delta(C_i)$ is the unit simplex in $\Re^{|C_i|}$. Note that $\rho \in \Delta(C_0)$. $\prod_{i\in N} \Delta(C_i)$ is the set of *mixed strategy profiles.* Given a mixed strategy profile $\mu = (\mu_i)_{i\in N} \in \prod_{i\in N} \Delta(C_i)$, define $\bar{\mu} = \prod_{i\in N-\{0\}} \mu_i \in \Delta(C_{-0})$. The probability measure generated on the measurable space $(C, 2^C)$ by the mixed strategy profile μ is Q_μ, defined by[4]

$$Q_\mu(E) = \begin{cases} \rho \times \bar{\mu}(E), & \text{if } 0 \in N \\ \bar{\mu}(E), & \text{if } 0 \notin N \end{cases}$$

for $E \in 2^C$. Note the use of the non-cooperative hypothesis in that the players randomize independently, hence the Cartesian product measure on C.

Given the probability space $(C, 2^C, Q_\mu)$, the distribution of $z(.)$ is $Q_\mu \circ z^{-1} \in \Delta(Z)$. It is easy to check that the play generated by $c \in C$ includes $t \in T$ iff. $z(c) \in Z_t$. Thus, the prior probability of node $t \in T$ being reached, when the mixed strategy profile μ is implemented, is

$$\begin{aligned} P_\mu(t) &= \begin{cases} \rho \times \bar{\mu}(\{c \in C \mid z(c) \in Z_t\}), & \text{if } 0 \in N \\ \bar{\mu}(\{c \in C \mid z(c) \in Z_t\}), & \text{if } 0 \notin N \end{cases} \\ &= \begin{cases} (\rho \times \bar{\mu}) \circ z^{-1}(Z_t), & \text{if } 0 \in N \\ \bar{\mu} \circ z^{-1}(Z_t), & \text{if } 0 \notin N \end{cases} \\ &= Q_\mu \circ z^{-1}(Z_t) \end{aligned}$$

It is easy to check that $P_\mu(t \vee t') = Q_\mu \circ z^{-1}(Z_t \cup Z_{t'})$ and $P_\mu(t \wedge t') = Q_\mu \circ z^{-1}(Z_t \cap Z_{t'})$; $P_\mu(t \vee t')$ is the prior probability of either node t or node t' being reached, and $P_\mu(t \wedge t')$ is the prior probability of node t and node t' being reached, when μ is implemented. Given $c \in C$, we write P_c instead of P_{δ_c}. The prior probability of information set s being reached by the play of the game, when the mixed strategy profile μ is implemented, is $P_\mu(s) = P_\mu(\vee_{t\in s} t) = Q_\mu \circ z^{-1}(\cup_{t\in s} Z_t)$. As the sets $\{Z_t \mid t \in s\}$ are mutually disjoint, $P_\mu(s) = \sum_{t\in s} Q_\mu \circ z^{-1}(Z_t) = \sum_{t\in s} P_\mu(t)$. Given these prior probabilities, we can use Bayes rule (when possible) to compute various conditional probabilities.

[4] As ρ is part of the description of the game, we shall lighten notation whenever possible by not mentioning it explicitly.

Given $i \in N - \{0\}$, define $u_i : C_{-0} \to \Re$ by

$$u_i(y) = \begin{cases} \sum_{x \in C_0} \rho(x) v_i \circ z(x,y), & \text{if } 0 \in N \\ v_i \circ z(y), & \text{if } 0 \notin N \end{cases}$$

for $y \in C_{-0}$.

Definition 1.2.1. *Given a standard $\Gamma(e)$, the strategic form of $\Gamma(e)$ is the collection $\Gamma = \big\{N - \{0\}, (C_i, u_i)_{i \in N - \{0\}}\big\}$.*

Thus, by definition, player 0 (Nature) is not involved in the strategic form description of a game. For $i \in N - \{0\}$, define $U_i : \prod_{i \in N} \Delta(C_i) \to \Re$ by

$$U_i(\mu) = \sum_{y \in C_{-0}} \bar{\mu}(y) u_i(y)$$

for $\mu \in \prod_{i \in N} \Delta(C_i)$. U_i is the expected utility function with u_i as the von Neumann-Morgenstern utility function. We now describe an alternative way of calculating U_i that will play a crucial role in Sections 1.3 and 1.4.

If $t \in Z$, then $P_\mu(t) = Q_\mu \circ z^{-1}(Z_t) = Q_\mu \circ z^{-1}(\{t\})$. Therefore, we have

$$\begin{aligned} \sum_{y \in C_{-0}} \bar{\mu}(y) u_i(y) &= \begin{cases} \sum_{(x,y) \in C_0 \times C_{-0}} \rho \times \bar{\mu}(x,y) v_i \circ z(x,y), & \text{if } 0 \in N \\ \sum_{y \in C_{-0}} \bar{\mu}(y) v_i \circ z(y), & \text{if } 0 \notin N \end{cases} \\ &= \sum_{c \in C} Q_\mu(c) v_i \circ z(c) \\ &= \sum_{t \in Z} Q_\mu \circ z^{-1}(\{t\}) v_i(t) \\ &= \sum_{t \in Z} Q_\mu \circ z^{-1}(Z_t) v_i(t) \\ &= \sum_{t \in Z} P_\mu(t) v_i(t) \end{aligned}$$

Thus, $U_i(\mu) = \sum_{t \in Z} P_\mu(t) v_i(t)$ for every $\mu \in \prod_{i \in N} \Delta(C_i)$.

Definition 1.2.3. *Given a strategic form game $\Gamma = \{N, (C_i, u_i)_{i \in N}\}$, the collection $\Gamma(m) = \{N, (\Delta(C_i), U_i)_{i \in N}\}$ is called the mixed extension of Γ.*

Given the expected utility form of U_i, a mixed strategy δ_c is equivalent to the pure strategy $c \in C_i$ as far as the decision-making of player i is concerned. If we choose some other extension of u_i, then this equivalence may not hold.

There is an alternative way of viewing the strategic choices faced by the players in a game $\Gamma(e)$. Every $i \in N$ is represented by a collection of *agents*, one for each $s \in S_i$, i.e., the set of agents of player i can be identified with S_i. A pure strategy of i can be interpreted as a set of instructions given by the player to his agents regarding the action they should take at their information set if that information set is reached by the play of the game. An artificial (but very useful) version of $\Gamma(e)$ is to consider all the agents $s \in S_i$ as distinct players. Every $s \in S_i$ is required to have the same information as i when called upon to move and the payoffs of s over the terminal nodes coincide with those of i. Thus, in terms of information and incentives, a player i and his agents are identical.

S is the set of players in the new game. The assignment of decision nodes $t \in X$ to this new set of players S is given by the mapping $t \mapsto s(t)$. The set of nodes controlled by player s is s. Player s controls only one information set, namely s. We define the utility assignment $\hat{v} : (S - S_0) \times Z \to \Re$ by $\hat{v}(s,t) = v(\iota(s), t)$ for every $s \in S - S_0$ and $t \in Z$; equivalently, $\hat{v}_s(t) = v_{\iota(s)}(t)$ for every $s \in S - S_0$ and $t \in Z$. These constructions yield a new extensive form game

$$\hat{\Gamma}(e) = \{(T, \succ), S, \iota, (D_t)_{t \in X}, S, \hat{v}, \rho\}$$

The strategic form of $\hat{\Gamma}(e)$, denoted by $\hat{\Gamma}$, is called the *agent normal form* of $\Gamma(e)$. The set of pure strategies of player $s \in S$ in $\hat{\Gamma}(e)$ is D_s. The set of mixed strategies of player $s \in S$ is $\Delta(D_s)$. The mixed extension of $\hat{\Gamma}$ is denoted by $\hat{\Gamma}^m$. In $\hat{\Gamma}(m)$, the consequence of agents $s \in S_i$ randomizing over their available actions is that the induced distribution over $\prod_{s \in S_i} D_s$ is a Cartesian product measure. This follows from the non-cooperative hypothesis that players randomize independently. In the case of $\Gamma(m)$, $\mu \in \Delta(C_i)$ does not necessarily induce a product measure over $\prod_{s \in S_i} D_s$; μ implicitly *correlates* the randomizations of a player at different information sets. We now define a type of strategy that is of great practical importance.

Definition 1.2.4. *Given a standard $\Gamma(e)$, the set of behavior strategies for player $i \in N$ is*

$$B_i = \begin{cases} \{\rho\}, & \text{if } i = 0 \\ \prod_{s \in S_i} \Delta(D_s), & \text{if } i \in N - \{0\} \end{cases}$$

The set of behavior strategy profiles is $B = \prod_{i \in N} B_i$.

If the players employ the behavior strategy profile $b \in B$, the prior probability of the play of the game reaching a node t, denoted by $P_b(t)$, is computed as follows. Let $\{t_0, \ldots, t_n\}$, with $t_n = t$, be the unique path from t_0 to t. Then,

$$P_b(t) = \prod_{i=0}^{n-1} b_{\iota(t_i)}(s(t_i)) \circ \alpha_{t_i}^{-1}(t_{i+1})$$

Define the function $V_i : B \to \Re$ by

$$V_i(b) = \sum_{t \in Z} P_b(t) v_i(t)$$

for $b \in B$ and $i \in N - \{0\}$; $V_i(b)$ is the expected payoff of player i when the profile b of behavior strategies is implemented. A behavior strategy $b \in B_i$ for player i generates a profile of mixed strategies $(\nu_s)_{s \in S_i} \in \prod_{s \in S_i} \Delta(D_s)$ for player i's agents, defined by

$$\nu_s = b(s) \tag{1.2.5}$$

for $s \in S_i$. Conversely, given a mixed strategy profile $(\nu_s)_{s \in S_i} \in \prod_{s \in S_i} \Delta(D_s)$ for player i's agents, (1.2.5) defines a behavior strategy $b \in B_i$ for player i. More generally, a behavior strategy profile $b = (b_i)_{i \in N} \in B$ generates a profile of mixed strategies $(\nu_s)_{s \in S} \in \prod_{s \in S} \Delta(D_s)$ for the agent normal form $\hat{\Gamma}$ by the formula

$$\nu_s = b_{\iota(s)}(s) \tag{1.2.6}$$

for $s \in S$. Conversely, given a profile of mixed strategies $(\nu_s)_{s \in S} \in \prod_{s \in S} \Delta(D_s)$ for the agent normal form $\hat{\Gamma}$, (1.2.6) generates a behavior strategy profile $b \in B$. Given a profile of mixed strategies $\nu = (\nu_s)_{s \in S}$ for the agent normal form $\hat{\Gamma}$, the prior probability of the play of the game reaching a node t, denoted by $P_\nu(t)$, is computed as follows. Let $\{t_0, \ldots, t_n\}$, with $t_n = t$, be the unique path from t_0 to t. Then,

$$P_\nu(t) = \prod_{i=0}^{n-1} \nu_{s(t_i)} \circ \alpha_{t_i}^{-1}(t_{i+1})$$

If the profile of behavior strategies, $b = (b_i)_{i \in N} \in B$, and the profile of mixed strategies for $\hat{\Gamma}$, $\nu = (\nu_s)_{s \in S} \in \prod_{s \in S} \Delta(D_s)$, satisfy (1.2.6), then

$$P_b(t) = \prod_{i=0}^{n-1} b_{\iota(t_i)}(s(t_i)) \circ \alpha_{t_i}^{-1}(t_{i+1}) = \prod_{i=0}^{n-1} \nu_{s(t_i)} \circ \alpha_{t_i}^{-1}(t_{i+1}) = P_\nu(t) \tag{1.2.7}$$

Consequently, if $\nu \in \prod_{s\in S}\Delta(D_s)$ and $b \in B$ satisfy (1.2.6), then for every $s \in S$,

$$V_{\iota(s)}(b) = \sum_{t\in Z} P_b(t) v_{\iota(s)}(t) = \sum_{t\in Z} P_\nu(t) v_{\iota(s)}(t) = \sum_{t\in Z} P_\nu(t)\hat{v}_s(t) \tag{1.2.8}$$

Note carefully the distinction between mixed and behavior strategies. A mixed strategy for player i calls for randomization over i's complete plans of action; a plan is (randomly) picked before the play of the game begins and decisions are made in accordance with the selected plan throughout the play of the game. On the other hand, a behavior strategy involves picking a plan of randomization over available actions at each of i's information sets; while a particular randomization plan is picked before the play of the game begins, the decisions at various information sets are not deterministically fixed by it. Thus, a mixed strategy can be interpreted as involving global randomization, while a behavior strategy is more appropriately interpreted as a sequence of local randomizations.

Definition 1.2.9. *Given a standard* $\Gamma(e)$, *the collection* $\Gamma(b) = \{N-\{0\}, (B_i, V_i)_{i\in N-\{0\}}\}$ *is called the behavior strategy extension of* $\Gamma(e)$.

1.3 A la Recherche du Temps Perdu

In Section 1.4 we shall define the notions of Nash equilibria in mixed strategies and Nash equilibria in behavior strategies, and consider the relationship between these notions. The crucial property required for this relationship to hold is a notion of equivalence between mixed strategies and behavior strategies that we define below. In this section we consider a condition on the extensive form game, called perfect recall, that is necessary and sufficient for the required equivalence to hold. We begin with some preliminary definitions.

Definition 1.3.1. *Consider a standard* $\Gamma(e)$.

(a) A node $t \in T$ *is said to be possible given* $c' \in C_i$ *if there exists a pure strategies profile* $c \in C$ *such that* $c_i = c'$ *and* $P_c(t) = 1$. $\mathrm{Poss}\,(c')$ *denotes the set of nodes that are possible given* c'.

(b) An information set s *is said to be relevant given* $c' \in C_i$ *if* $s \cap \mathrm{Poss}\,(c') \neq \emptyset$. $\mathrm{Rel}\,(c')$ *denotes the collection of information sets that are relevant given* c'.

(c) An information set s *is said to be relevant given* $\mu \in \Delta(C_i)$ *if there exists a pure*

strategy $c' \in \operatorname{supp}\mu$ *such that* $s \in \operatorname{Rel}(c')$. $\operatorname{Rel}(\mu)$ *denotes the collection of information sets that are relevant given* μ.

Given a pure strategy $c' \in C_i$, $\operatorname{Poss}(c')$ is the set of nodes that can *possibly* be reached by the play of the game under *some* profile of pure strategies $c_{-i} \in C_{-i}$ for the other players. $T - \operatorname{Poss}(c')$ is the set of nodes that cannot possibly be reached by the play of the game if player i implements c'. Analogously, given a pure strategy $c' \in C_i$, $\operatorname{Rel}(c')$ is the collection of information sets that can *possibly* be reached by the play of the game under *some* profile of pure strategies $c_{-i} \in C_{-i}$ for the other players. $S - \operatorname{Rel}(c')$ is the collection of information sets that cannot possibly be reached by the play of the game if player i implements c'. Note that, given a mixed strategy $\mu \in \Delta(C_i)$, $\operatorname{Rel}(\mu) = \cup_{c' \in \operatorname{supp}\mu} \operatorname{Rel}(c')$.

Consider a standard game $\Gamma(e)$ and $i \in N$. The *mixed strategy representation* $\mu \in \Delta(C_i)$ of the behavior strategy $b \in B_i$ is given by

$$\mu(c) = \prod_{s \in S_i} b(s)\big(c(s)\big) \tag{1.3.2}$$

for $c \in C_i$. It is easy to verify, using an argument similar to that used in the proof of Lemma 1.3.4(a), that μ defined by (1.3.2) is indeed a mixed strategy for i. $\mu \in \prod_{i \in N} \Delta(C_i)$ is called the mixed strategy representation of the behavior strategy profile $b \in B$ if for every $i \in N$, μ_i is the mixed strategy representation of b_i.

Given an information set $s \in S_i$ and decision $d \in D_s$, let $C_i(s) = \{c \in C_i \mid s \in \operatorname{Rel}(c)\}$ and $C_i(s,d) = \{c \in C_i \mid c(s) = d\}$. The *behavior strategy representation* $b \in B_i$ of a mixed strategy $\mu \in \Delta(C_i)$ is given by

$$b(s)(d) = \begin{cases} \mu(C_i(s) \cap C_i(s,d))/\mu(C_i(s)), & \text{if } \mu(C_i(s)) > 0 \\ \mu(C_i(s,d)), & \text{if } \mu(C_i(s)) = 0 \end{cases} \tag{1.3.3}$$

for $s \in S_i$ and $d \in D_s$. The profile $b \in B$ is called the behavior strategy representation of the mixed strategy profile $\mu \in \prod_{i \in N} \Delta(C_i)$ if for every $i \in N$, b_i is the behavior strategy representation of μ_i.

Lemma 1.3.4. *Consider a standard* $\Gamma(e)$.

(A) Let $\mu \in \Delta(C_i)$ *be the mixed representation of* $b \in B_i$. *If* b^* *is the behavior representation of* μ, *then* $b = b^*$.

(B) Let $b \in B_i$ be the behavior representation of $\mu \in \Delta(C_i)$. If μ^* is the mixed representation of b, then μ^* and μ may not coincide.

Proof. (A) Let $S_i = \{s_1, \ldots, s_n\}$. We show that $b(s_1) = b^*(s_1)$; the argument for other information sets is analogous. Let $d^* \in D_{s_1}$. As $C_i(s_1)$ and $C_i(s_1, d^*)$ are independent, it follows that $\mu(C_i(s_1) \cap C_i(s_1, d^*)) = \mu(C_i(s_1)).\mu(C_i(s_1, d^*))$. Using the definition of C_i, it follows that $C_i(s_1, d^*) = \{d \in \prod_{j=1}^n D_{s_j} \mid d_1 = d^*\}$. Thus, by (1.3.2)

$$\mu(C_i(s_1, d^*)) = \sum_{c \in C_i(s_1, d^*)} \mu(c) = \sum_{c \in C_i(s_1, d^*)} \prod_{j=1}^n b(s_j)(c(s_j))$$

As $c(s_1) = d^*$ for every $c \in C_i(s_1, d^*)$, we have

$$\begin{aligned}
\sum_{c \in C_i(s_1, d^*)} \prod_{j=1}^n b(s_j)(c(s_j)) &= b(s_1)(d^*) \sum_{c \in C_i(s_1, d^*)} \prod_{j=2}^n b(s_j)(c(s_j)) \\
&= b(s_1)(d^*) \sum_{d_{-1} \in \prod_{j=2}^n D_{s_j}} \prod_{j=2}^n b(s_j)(d_j)
\end{aligned}$$

Now

$$\begin{aligned}
\sum_{d_{-1} \in \prod_{j=2}^n D_{s_j}} \prod_{j=2}^n b(s_j)(d_j) &= \sum_{d_2 \in D_{s_2}} \cdots \sum_{d_n \in D_{s_n}} \prod_{j=2}^n b(s_j)(d_j) \\
&= \sum_{d_2 \in D_{s_2}} \cdots \sum_{d_{n-1} \in D_{s_{n-1}}} \prod_{j=2}^{n-1} b(s_j)(d_j) \sum_{d_n \in D_{s_n}} b(s_n)(d_n) \\
&= \sum_{d_2 \in D_{s_2}} \cdots \sum_{d_{n-1} \in D_{s_{n-1}}} \prod_{j=2}^{n-1} b(s_j)(d_j)
\end{aligned}$$

Repeating this step $n-2$ times, we have

$$\sum_{d_{-1} \in \prod_{j=2}^n D_{s_j}} \prod_{j=2}^n b(s_j)(d_j) = 1$$

Therefore, $\mu(C_i(s_1, d^*)) = b(s_1)(d^*)$ and $\mu(C_i(s_1) \cap C_i(s_1, d^*)) = \mu(C_i(s_1)).b(s_1)(d^*)$. If $\mu(C_i(s_1)) > 0$, then (1.3.3) implies

$$b(s_1)(d^*) = \frac{\mu(C_i(s_1) \cap C_i(s_1, d^*))}{\mu(C_i(s_1))} = b^*(s_1)(d^*)$$

If $\mu(C_i(s_1)) = 0$, then (1.3.3) implies $b^*(s_1)(d^*) = \mu(C_i(s_1, d^*)) = b(s_1)(d^*)$.

(B) can be seen from a simple example. ■

We now define the notion of equivalence that we alluded to at the beginning of this section.

Definition 1.3.5. *Consider a standard* $\Gamma(e)$*. Let* $b \in B$ *be the behavior representation of the mixed strategy profile* $\mu \in \prod_{i \in N} \Delta(C_i)$*. We say that* μ *and* b *are equivalent if* $P_\mu(t) = P_b(t)$ *for every* $t \in Z$*.*

Such an equivalence is of great importance in the analysis of extensive form games because it ensures that expected payoff calculations are unaffected by our choice between mixed strategy profiles and their behavior strategy representations. Kuhn's theorem provides a very general (and natural) sufficient condition, called *perfect recall*, for this equivalence to hold.

Definition 1.3.6. *(Kuhn) A standard game* $\Gamma(e)$ *is said to satisfy perfect recall if for every* $i \in N$*,* $s \in S_i$ *and* $c \in C_i$*,*

$$\langle s \in \mathrm{Rel}\,(c) \rangle \qquad \Longrightarrow \qquad \langle s \subset \mathrm{Poss}\,(c) \rangle$$

The following is the equivalence theorem and a corollary.

Theorem 1.3.7. *(Kuhn) Consider a standard* $\Gamma(e)$*. Suppose* $b \in B$ *is the behavior representation of the mixed strategy profile* $\mu \in \prod_{i \in N} \Delta(C_i)$*. If* $\Gamma(e)$ *satisfies perfect recall, then* μ *and* b *are equivalent. Conversely, if the mixed strategy profile* μ *and its behavior representation* b *are equivalent for every* $\mu \in \prod_{i \in N} \Delta(C_i)$*, then* $\Gamma(e)$ *satisfies perfect recall.*

Corollary 1.3.8. *Consider a standard* $\Gamma(e)$*. Suppose* b *is the behavior representation of the mixed strategy profile* μ*. If* $\Gamma(e)$ *satisfies perfect recall, then* $U_i(\mu) = V_i(b)$ *for every* $i \in N - \{0\}$*.*

1.4 Nash equilibrium

The above sections provide a formal description of finite games. The main question that game theory attempts to answer is: given a game, what are the likely outcomes[5] of the game? Answering this question involves formulating a general theory of how these predictions should be made. Merely predicting outcomes is not enough. We wish to provide

[5] This is deliberately vague. Depending on our needs, an "outcome" can refer to either a terminal node, a play of the game or a profile of strategies.

general principles that explain why certain predictions in some class of games should be considered compelling while others are not so. An *equilibrium (or solution) concept* is the taxonomical tool used to separate out the set of compelling predictions from the rest. More formally, an equilibrium concept is a mapping defined over a class of games that maps any given game from that class into a set of strategy profiles that are interpreted to be compelling predictions of the outcome of the game. The main solution concept used in non-cooperative game theory is that of *Nash equilibrium*.[6]

Definition 1.4.1. *Given a strategic form game $\Gamma = \{N, (C_i, u_i)_{i \in N}\}$, a profile $c \in C$ is called a pure strategy Nash equilibrium if, for every $i \in N$ and $x \in C_i$, $u_i(c) \geq u_i(c_{-i}, x)$.*

An important problem with this notion is that many games do not have a pure strategy Nash equilibrium; i.e., the concept may be vacuous for many games. The idea of a mixed extension was developed precisely to overcome this problem. It turns out that the mixed extension of a finite game always has a Nash equilibrium; this will be proved in the next section.

Definition 1.4.2. *Consider the strategic form game $\Gamma = \{N, (C_i, u_i)_{i \in N}\}$. Let $\Gamma(m)$ be the mixed extension of Γ. A profile of mixed strategies $\mu \in \prod_{i \in N} \Delta(C_i)$ is called a mixed strategy Nash equilibrium of Γ if for every $i \in N$ and $\lambda \in \Delta(C_i)$, $U_i(\mu) \geq U_i(\mu_{-i}, \lambda)$.*

Since pure strategies are merely special types of mixed strategies, we shall use "equilibrium" to refer to a mixed strategy Nash equilibrium; when specifically referring to a pure strategy equilibrium, we append the qualifier "pure". Given an extensive form game $\Gamma(e)$, a *Nash equilibrium of* $\Gamma(e)$ refers to a Nash equilibrium of its strategic form representation Γ. However, even in the simplest extensive form games, calculation of mixed strategy equilibria can be quite laborious. This becomes even more difficult with complex theoretical models in which one might want simple transparent characterizations of the equilibria. We now consider an alternative way of calculating and characterizing mixed strategy Nash equilibria of $\Gamma(e)$.

Consider the agent normal form $\hat{\Gamma}$ of the game. Any equilibrium profile $(\nu_s)_{s \in S}$ of the agent normal form generates, via (1.2.5), a profile of behavior strategies $(b_i)_{i \in N}$. We can

[6] See Nash (1950) for the formulation and existence proof.

interpret these behavior strategies as representing "equilibrium behavior". The problem is that the mixed representations of these constructed behavior strategies often do *not* constitute a Nash equilibrium of the strategic form of $\Gamma(e)$. The reason for this problem is clear. While the agent normal form preserves the information and preference structure of $\Gamma(e)$, it rules out correlation of actions across different information sets controlled by the same player. Given the non-cooperative nature of the game, the agents of a given player cannot correlate their actions, while a mixed strategy in the strategic form of the game allows such correlation.

Definition 1.4.3. *Given a standard $\Gamma(e)$, a profile of behavior strategies $b = (b_i)_{i \in N - \{0\}} \in B_{-0}$ is called a Nash equilibrium in behavior strategies if*

(a) $\nu = (\nu_s)_{s \in S - S_0} = \big(b_{\iota(s)}(s)\big)_{s \in S - S_0}$ is a Nash equilibrium of the agent normal form $\hat{\Gamma}$, and

(b) the mixed representation of b is a Nash equilibrium of Γ.

Given this definition, it does not appear to be any easier to compute a Nash equilibrium in behavior strategies than to calculate one in mixed strategies. However, the following characterization result greatly simplifies the task.

Theorem 1.4.4. *Given a standard $\Gamma(e)$ that satisfies perfect recall, $b = (b_i)_{i \in N - \{0\}} \in B_{-0}$ is a Nash equilibrium in behavior strategies iff. for every $i \in N - \{0\}$ and $b'_i \in B_i$*

$$V_i(b) \geq V_i(b_{-i}, b'_i) \tag{1.4.5}$$

This shows that for games with perfect recall, we may restrict our search for equilibria to Nash equilibria of the behavior extension $\Gamma(b)$ of the given game $\Gamma(e)$, i.e., the game in which the players are restricted to choosing from among their behavior strategies. We defer the proof of this result till after we have proved some elementary lemmas.

Lemma 1.4.6. *Let $\hat{\Gamma}$ be the agent normal form of a standard $\Gamma(e)$. If a profile $b \in B_{-0}$ satisfies (1.4.5), then the profile $\nu = (\nu_s)_{s \in S - S_0}$, defined by (1.2.5), is a Nash equilibrium of $\hat{\Gamma}$.*

Proof. Suppose $b \in B_{-0}$ satisfies (1.4.5) for every i, ν is derived from b by (1.2.6), and ν is not a Nash equilibrium of $\hat{\Gamma}$. Then there exists $s \in S - S_0$ and $\nu'_s \in \Delta(D_s)$ such that

$\sum_{t\in Z} P_{\nu_{-s},\nu'_s}(t)\hat{v}_s(t) > \sum_{t\in Z} P_\nu(t)\hat{v}_s(t)$. Define, for $s' \in S - S_0$,

$$b'_{\iota(s)}(s') = \begin{cases} \nu_{s'}, & \text{if } s' \neq s \\ \nu'_s, & \text{if } s' = s \end{cases}$$

Then, using (1.2.8),

$$V_{\iota(s)}(b_{-\iota(s)}, b'_{\iota(s)}) = \sum_{t\in Z} P_{\nu_{-s},\nu'_s}(t)\hat{v}_s(t) > \sum_{t\in Z} P_\nu(t)\hat{v}_s(t) = V_{\iota(s)}(b)$$

which yields a contradiction. ■

Lemma 1.4.7. *Consider a standard $\Gamma(e)$ that satisfies perfect recall.*

(A) If μ^1 and μ^2 are profiles of mixed strategies with the same behavior representation b, and μ^1 is a Nash equilibrium, then μ^2 also is a Nash equilibrium.

(B) A mixed representation μ of a behavior strategy profile b is a Nash equilibrium iff. every mixed strategy profile μ' that has b as the behavior representation is a Nash equilibrium.

(C) Let b be the behavior representation of a mixed strategy profile μ. If μ is a Nash equilibrium of Γ, then b satisfies (1.4.5).

Proof. (A) Suppose μ^2 is not a Nash equilibrium. Then there exists $i \in N - \{0\}$ and $\mu_i \in \Delta(C_i)$ such that $U_i(\mu^2_{-i}, \mu_i) > U_i(\mu^2)$. Let b'_i be the behavior representation of μ_i. Using Corollary 1.3.8, we have $V_i(b_{-i}, b'_i) = U_i(\mu^2_{-i}, \mu_i) > U_i(\mu^2) = V_i(b)$. Let μ'_i be the mixed representation of b'_i. By Lemma 1.3.4(A), (b_{-i}, b'_i) is the behavior representation of (μ^1_{-i}, μ'_i). Using Corollary 1.3.8, we have $U_i(\mu^1_{-i}, \mu'_i) = V_i(b_{-i}, b'_i) > V_i(b) = U_i(\mu^1)$, which is a contradiction.

(B) Suppose μ is the mixed representation of the behavior strategy profile b and μ is a Nash equilibrium. Then, by Lemma 1.3.4(A), b is the behavior representation of μ. Let μ' be another profile with b as the behavior representation. By (A), μ' also is a Nash equilibrium. Conversely, suppose every mixed strategy profile μ with behavior representation b is a Nash equilibrium. If μ' is a mixed representation of b, then b is the behavior representation of μ' by Lemma 1.3.4(A), and therefore μ' is a Nash equilibrium.

(C) Suppose b does not satisfy (1.4.5). Then there exists $i \in N$ and $b'_i \in B_i$ such that $V_i(b_{-i}, b'_i) > V_i(b)$. Let μ'_i be the mixed representation of b'_i. Then, by Lemma 1.3.4(A), (b_{-i}, b'_i) is the behavior representation of (μ_{-i}, μ'_i). Using Corollary 1.3.8, $U_i(\mu_{-i}, \mu'_i) = V_i(b_{-i}, b'_i) > V_i(b) = U_i(\mu)$, which is a contradiction. ■

Proof of Theorem 1.4.4. Suppose b is a profile of behavior strategies satisfying (1.4.5). Let μ be the mixed representation of b. Suppose μ is not a Nash equilibrium. Then there exists $i \in N$ and $\mu'_i \in \Delta(C_i)$ such that $U_i(\mu_{-i}, \mu'_i) > U_i(\mu)$. By Lemma 1.3.4(A), b is the behavior representation of μ. Let b'_i be the behavior representation of μ'_i. Using Corollary 1.3.8, $V_i(b_{-i}, b'_i) = U_i(\mu_{-i}, \mu'_i) > U_i(\mu) = V_i(b)$, which contradicts (1.4.5). It follows from Lemma 1.4.6 that if b satisfies (1.4.5), then 1.4.3(a) is satisfied.

Conversely, suppose b is a Nash equilibrium in behavior strategies. By definition, the mixed representation μ of b is a Nash equilibrium. By Lemma 1.3.4(A), b is the behavior representation of μ. Therefore, by Lemma 1.4.7(C), b must satisfy (1.4.5). ■

1.5 Existence of Nash equilibrium

In this section we establish sufficient conditions under which Nash equilibria exist. The essential steps in the proofs are provided below. The details can be filled in by the interested reader.

Theorem 1.5.1. *(Nash) If the collection* $\Gamma(a) = \{N, (\Sigma_i, u_i)_{i \in N}\}$ *is such that*

(a) N *is nonempty and finite, and*

for every $i \in N$,

(b) Σ_i *is a nonempty, compact and convex subset of* $\Re^l$, *where* $l \in \mathcal{N}$,

(c) $u_i : \Sigma \to \Re$ *is continuous, and*

(d) for every $\sigma_{-i} \in \Sigma_{-i}$, $u_i(\sigma_{-i}, .) : \Sigma_i \to \Re$ *is quasi-concave,*

then there exists $\sigma \in \Sigma$ *such that, for every* $i \in N$ *and* $\sigma' \in \Sigma_i$, $u_i(\sigma) \geq u_i(\sigma_{-i}, \sigma')$.

We can interpret $\Gamma(a)$ as an abstract game, i.e., one whose extensive form is not given. In this interpretation, N is the set of players, Σ_i is the set of strategies available to player i, and u_i is player i's payoff function. Note that an abstract game is formally analogous to a strategic form game.

In the above theorem, the profile σ that satisfies the stated inequalities is called a Nash equilibrium. An application of this theorem provides a general set of sufficient conditions for any finite strategic form game to possess a Nash equilibrium in mixed strategies. Let Γ be a finite strategic form game. If we define Σ_i to be the set of pure strategies C_i, then the game need not have a Nash equilibrium (in pure strategies). The reasons are evident in the light of Nash's theorem. C_i is nonempty and compact, and u_i is continuous. However, in the context of pure strategies, condition (d) is meaningless since convex combinations of strategies are not well-defined, i.e., C_i is not a convex set. The solution to this problem is obvious. Let $\Gamma(m)$ be the mixed extension of Γ. Now set $\Sigma_i = \Delta(C_i)$, and extend u_i to the expected utility function U_i. It is obvious that the collection $\Gamma(m)$ is an abstract game. Therefore, we have

Corollary 1.5.2. *If Γ is a finite strategic form game, then the mixed extension $\Gamma(m)$ has a Nash equilibrium, i.e., Γ has a mixed strategy Nash equilibrium.*

We begin the proof of Theorem 1.5.1 with the following definition.

Definition 1.5.3. *Given $\Gamma(a)$, the best-response mapping of player i is $R_i : \Sigma \Rightarrow \Sigma_i$, defined by*

$$R_i(\sigma) = \bigcap_{\sigma' \in \Sigma_i} \{\hat{\sigma} \in \Sigma_i \mid u_i(\sigma_{-i}, \hat{\sigma}) \geq u_i(\sigma_{-i}, \sigma')\}$$

Proof of Theorem 1.5.1.[7] By Berge's theorem of the maximum, R_i has nonempty values and $\mathrm{Gr}\, R_i$ is closed in $\Sigma \times \Sigma_i$ for every $i \in N$. Furthermore, R_i has convex values for every $i \in N$. Define $R : \Sigma \Rightarrow \Sigma$ by $R(\sigma) = \prod_{i \in N} R_i(\sigma)$. It can be verified that R has nonempty and convex values, and $\mathrm{Gr}\, R$ is closed in $\Sigma \times \Sigma$. It follows from Kakutani's fixed-point theorem that there exists $\sigma \in \Sigma$ such that $\sigma \in R(\sigma)$. This means $\sigma_i \in R_i(\sigma)$ for every $i \in N$, which concludes the proof. ■

While this result is adequate for proving the existence of (mixed strategy) Nash equilibria for finite games (Corollary 1.5.2), more powerful results are needed for games in which the strategy spaces are not finite-dimensional. Here is one of the many generalizations of Nash's original theorem.

Theorem 1.5.5. *(Browder) If the collection $\Gamma(a) = \{N, (\Sigma_i, u_i)_{i \in N}\}$ is such that,*

(a) N is nonempty and finite, and

for every $i \in N$,

(b) Σ_i is a nonempty, compact and convex subset of a Hausdorff real linear topological space X_i,

[7] This is a sketch of the proof. The important details to be taken care of are regarding the properties of product sets and product mappings. Assumption of the product topology and use of Tychonoff's theorem take care of these details. The same comments apply to the proof of Theorem 1.5.8.

(c) $u_i : \Sigma \to \Re$ *is continuous, and*

(d) for every $\sigma_{-i} \in \Sigma_{-i}$, $u_i(\sigma_{-i}, .) : \Sigma_i \to \Re$ *is quasi-concave,*

then there exists $\sigma \in \Sigma$ *such that, for every* $i \in N$ *and* $\sigma' \in \Sigma_i$, $u_i(\sigma) \geq u_i(\sigma_{-i}, \sigma')$.

Among other things, this theorem allows us to extend Corollary 1.5.2 very substantially.

Theorem 1.5.6. *(Glicksberg) If the collection* $\Gamma = \{N, (\Sigma_i, u_i)_{i \in N}\}$ *is such that,*

(a) N *is nonempty and finite, and*

for every $i \in N$

(b) Σ_i *is a nonempty, compact metric space, and*

(c) $u_i : \Sigma \to \Re$ *is continuous,*

then there exists $\mu \in \prod_{i \in N} \Delta(\Sigma_i)$ *such that, for every* $i \in N$ *and* $\mu' \in \Delta(\Sigma_i)$,

$$\int_\Sigma \left(\prod_{j \in N} \mu_j \right) (d\sigma) u_i(\sigma) \geq \int_\Sigma \left(\mu' \times \prod_{j \in N - \{i\}} \mu_j \right) (d\sigma) u_i(\sigma)$$

Numerous other such results are available in a variety of contexts; see Debreu [], Ky Fan [], Mas-Colell [], Milgrom and Weber [].

Consider a collection of identical sets $\{\Sigma_i \mid i \in N\}$ and let $\Sigma = \prod_{i \in N} \Sigma_i$. Define $\operatorname{diag} \Sigma = \{\sigma \in \Sigma \mid \sigma_i = \sigma_j, \quad \forall i, j \in N\}$. A bijection $\pi : N \to N$ is called a *permutation* on N. Given a profile of strategies $\sigma \in \Sigma$ and a permutation $\pi : N \to N$, define the profile $\sigma^\pi \in \Sigma$ by: $\sigma_i^\pi = \sigma_{\pi(i)}$ for every $i \in N$.

Definition 1.5.7. *A game* $\Gamma(a) = \{N, (\Sigma_i, u_i)_{i \in N}\}$ *is said to be symmetric if*

(a) $\Sigma_i = \Sigma_j$ *for all* $i, j \in N$, *and*

(b) for every permutation $\pi : N \to N$ *and* $i \in N$, $u_i(\sigma^\pi) = u_{\pi(i)}(\sigma)$.

A very useful fact regarding symmetric games is that (under appropriate conditions) they have symmetric equilibria. We prove this result in the same setting as Theorem 1.5.1; analogous results can be derived for other settings.

Theorem 1.5.8. *If* $\Gamma(a) = \{N, (\Sigma_i, u_i)_{i \in N}\}$ *is a symmetric game such that,*

(a) N *is nonempty and finite, and*

for every $i \in N$

(b) Σ_i *is a nonempty, compact and convex subset of* $\Re^l$, *where* $l \in \mathcal{N}$,

(c) u_i *is continuous, and*

(d) for every $\sigma_{-i} \in \Sigma_{-i}$, $u_i(\sigma_{-i}, .) : \Sigma_i \to \Re$ *is quasi-concave,*

then there exists $\sigma \in \operatorname{diag}\Sigma$ *such that, for every* $i \in N$ *and* $\sigma' \in \Sigma_i$, $u_i(\sigma) \geq u_i(\sigma_{-i}, \sigma')$.

Proof. It is easy to check that $\operatorname{diag}\Sigma$ is nonempty, convex and compact (since it is a closed subset of Σ which is compact). Moreover, for every $\sigma \in \operatorname{diag}\Sigma$, $R(\sigma) \cap \operatorname{diag}\Sigma \neq \emptyset$. Define $\bar{R} : \operatorname{diag}\Sigma \Rightarrow \operatorname{diag}\Sigma$ by $\bar{R}(\sigma) = R(\sigma) \cap \operatorname{diag}\Sigma$. $\bar{R}$ has nonempty and convex values, and $\operatorname{Gr}\bar{R}$ is closed in $\operatorname{diag}\Sigma \times \operatorname{diag}\Sigma$. By Kakutani's fixed point theorem (Kakutani 1941), there exists $\sigma \in \operatorname{diag}\Sigma$ such that $\sigma \in \bar{R}(\sigma)$. By the definition of $\bar{R}$, $\sigma \in R(\sigma)$, which concludes the proof. ■

1.6 Two-person zero-sum games

We consider abstract games of the form $\Gamma(a) = \{\{1,2\}, (\Sigma_i, u_i)_{i=1,2}\}$, where $u_2 = -u_1$. Given this condition, it suffices to specify one payoff function only. Henceforth, for the purposes of this section, we shall denote u_1 by u. For zero-sum games, the Nash equilibria have a characterization which makes them particularly compelling as predictions of the outcome of the game. This characterization is contained in the famous minmax theorem. Before stating and proving this theorem, we note a simple lemma.

Lemma 1.6.1. *Consider an abstract zero-sum game* $\Gamma(a) = \{\{1,2\}, (\Sigma_1, \Sigma_2), u\}$. *Suppose*

(a) Σ_1 *and* Σ_2 *are nonempty and compact subsets of* $\Re^l$*, and*

(b) u *is continuous.*

Then,

(A) The functions $\bar{u} : \Sigma_2 \to \Re$ *and* $\underline{u} : \Sigma_1 \to \Re$*, defined by*

$$\bar{u}(\tau_2) = \max_{\tau_1 \in \Sigma_1} u(\tau_1, \tau_2) \qquad \text{and} \qquad \underline{u}(\tau_1) = \min_{\tau_2 \in \Sigma_2} u(\tau_1, \tau_2)$$

are well-defined and continuous.

(B) For every $(\tau_1, \tau_2) \in \Sigma$, $\underline{u}(\tau_1) \leq u(\tau_1, \tau_2)$ *and* $\bar{u}(\tau_2) \geq u(\tau_1, \tau_2)$.

(C) $\min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2) \geq \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1)$.

(D) If $(\sigma_1, \sigma_2) \in \Sigma$ *is a Nash equilibrium of* $\Gamma(a)$*, then*

$$u(\sigma_1, \sigma_2) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1) = \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2)$$

Proof. (A) Using Weierstrass' theorem, these functions are well-defined. Continuity follows from the theorem of the maximum (see Berge).

(B) Follows from the definitions.

(C) It follows from (B) that $\bar{u}(\tau_2) \geq u(\tau_1, \tau_2) \geq \underline{u}(\tau_1)$. Since this relation holds for every $\tau_1 \in \Sigma_1$ and $\tau_2 \in \Sigma_2$, the result follows.

(D) By (A) and Weierstrass' theorem, $\max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1)$ and $\min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2)$ are well-defined. Then,

$$u(\sigma_1, \sigma_2) = \bar{u}(\sigma_2) \geq \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2) \geq \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1) \geq \underline{u}(\sigma_1) = u(\sigma_1, \sigma_2)$$

The equalities follow from the fact that (σ_1, σ_2) is an equilibrium. The first and third inequalities follow from the definitions of maximum and minimum. The second inequality follows from (C). ■

In a zero-sum game, player 1's objective is to maximize u and player 2's objective is to minimize u. Given that player 1 plays $\tau_1 \in \Sigma_1$, $\underline{u}(\tau_1)$ is the lowest payoff to which player 2 can restrict player 1; equivalently, given $\tau_1 \in \Sigma_1$, $-\underline{u}(\tau_1)$ is the highest payoff that player 2 can get. Correspondingly, given that player 2 plays $\tau_2 \in \Sigma_2$, $\bar{u}(\tau_2)$ is the highest payoff that player 1 can get; equivalently, given $\tau_2 \in \Sigma_2$, $-\bar{u}(\tau_2)$ is the lowest payoff to which player 2 can be restricted by player 1.

Let $\sigma_1 \in \Sigma_1$ be such that $\underline{u}(\sigma_1) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1)$. By Lemma 1.6.1(B), $\underline{u}(\sigma_1) \leq u(\sigma_1, \tau_2)$ for every $\tau_2 \in \Sigma_2$. Thus, $\underline{u}(\sigma_1)$ is guaranteed, regardless of what player 2 does. As $\underline{u}(\tau_1)$ is the payoff that player 1 can guarantee by playing τ_1, and $\underline{u}(\sigma_1) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1)$, $\underline{u}(\sigma_1)$ is the highest guaranteed payoff for player 1, regardless of what player 2 does.

Theorem 1.6.2. *Consider an abstract game* $\Gamma(a) = \{\{1, 2\}, (\Sigma_1, \Sigma_2), u\}$. *Suppose*

(a) Σ_1 *and* Σ_2 *are nonempty, convex and compact subsets of* $\Re^l$,

(b) u *is continuous, and*

(c) $u(., \tau_2) : \Sigma_1 \to \Re$ *is quasi-concave for every* $\tau_2 \in \Sigma_2$ *and* $u(\tau_1, .) : \Sigma_2 \to \Re$ *is quasi-convex for every* $\tau_1 \in \Sigma_1$.

Then,

(A) $(\sigma_1, \sigma_2) \in \Sigma$ *is a Nash equilibrium of* $\Gamma(a)$ *iff.*

$$\underline{u}(\sigma_1) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1) \qquad \text{and} \qquad \bar{u}(\sigma_2) = \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2)$$

(B) If $(\sigma_1, \sigma_2) \in \Sigma$ and $(\sigma_1', \sigma_2') \in \Sigma$ are Nash equilibria of $\Gamma(a)$, then (σ_1, σ_2') and (σ_1', σ_2) are Nash equilibria of $\Gamma(a)$.

Proof. (A) Let (σ_1, σ_2) be a Nash equilibrium of $\Gamma(a)$. Note that

$$\underline{u}(\sigma_1) = u(\sigma_1, \sigma_2) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1) \geq \underline{u}(\tau_1)$$

for every $\tau_1 \in \Sigma_1$ and

$$\bar{u}(\sigma_2) = u(\sigma_1, \sigma_2) = \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2) \leq \bar{u}(\tau_2)$$

for every $\tau_2 \in \Sigma_2$. The first equalities follow from the definition of equilibrium. The second equalities follow from Lemma 1.6.1(D). The inequalities follow from the definitions of maximum and minimum. Conversely, let $\sigma_1 \in \Sigma_1$ and $\sigma_2 \in \Sigma_2$ be such that $\underline{u}(\sigma_1) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1)$ and $\bar{u}(\sigma_2) = \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2)$. Since $\Gamma(a)$ satisfies the hypotheses of Theorem 8.5.1, it must have a Nash equilibrium, say (σ_1^*, σ_2^*). By Lemma 1.6.1(D), the following relation must hold

$$u(\sigma_1^*, \sigma_2^*) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1) = \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2)$$

Using this relation, the given hypothesis, and Lemma 1.6.1(B), we have

$$u(\sigma_1, \sigma_2) \geq \underline{u}(\sigma_1) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1) = \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2) = \bar{u}(\sigma_2) \geq u(\sigma_1, \sigma_2)$$

This means $u(\sigma_1, \sigma_2) = \underline{u}(\sigma_1) = \bar{u}(\sigma_2)$, which proves the result.

(B) Since (σ_1, σ_2) is a Nash equilibrium, by (A) we must have $\underline{u}(\sigma_1) = \max_{\tau_1 \in \Sigma_1} \underline{u}(\tau_1)$. By the same argument, $\bar{u}(\sigma_2') = \min_{\tau_2 \in \Sigma_2} \bar{u}(\tau_2)$, since (σ_1', σ_2') is a Nash equilibrium. Therefore, it follows from (A) that (σ_1, σ_2') is a Nash equilibrium. ■

The above theorem contains some important properties of two player zero-sum games. First, for a given player, the payoff in every Nash equilibrium of the game is exactly the same. In every Nash equilibrium, the maximizing player gets the maxmin payoff (which is equal to the minmax payoff) and the minimizing player gets the negative of that payoff. Secondly, the strategies employed in different Nash equilibria are interchangeable. It is these two properties that make Nash equilibria especially compelling as predictions of outcomes in two player zero-sum games.

1.7 Subgame perfect Nash equilibrium

As many games have a large number of Nash equilibria, the Nash equilibrium concept does not sufficiently restrict the set of predictions. Therefore, many additional restrictions have been proposed to refine the set of Nash equilibria. These refinements usually impose stronger rationality conditions on player behavior than is required by Nash equilibrium. Like Nash equilibrium, some of the proposed refinements require only the data contained in the strategic form of a given game, while others use the information contained in the extensive form of the game. A leading example of the latter family of refinements is the notion of a subgame perfect Nash equilibrium.[8] Informally, a Nash equilibrium of a given game $\Gamma(e)$ is said to be subgame perfect if it "induces a Nash equilibrium in every proper subgame".

Definition 1.7.1. *Consider a standard game* $\Gamma(e)$ *with perfect recall and a Nash equilibrium* $b = (b_i)_{i \in N-\{0\}}$ *of* $\Gamma(e)$. b *is called a subgame perfect Nash equilibrium of* $\Gamma(e)$ *if, for every proper subgame of* $\Gamma(e)$, *the restriction of* b *to that subgame constitutes a Nash equilibrium for that subgame.*

This refinement serves to eliminate Nash equilibria that involve sub-optimal behavior by a player at some information set. Clearly, given a Nash equilibrium profile of strategies, such sub-optimal behavior cannot occur at any information set that is reached with positive probability by the play generated by the given profile. In many applications, such sub-optimal planned behavior may be interpreted as a non-credible threat, i.e., a threat that the relevant player will be irrational to carry out, should his bluff be called.

[8] See Selten (1975) for the formulation.

2. Information structures in games

2.1 Perfect information games

Definition 2.1.1. *An extensive form game* $\Gamma(e)$ *is said to be a perfect information game if every information set of* $\Gamma(e)$ *is a singleton.*

A trivial observation that is implied by this definition is the following.

Theorem 2.1.2. *Every perfect information game satisfies perfect recall.*

We now note a very useful property called the "single deviation property". We simplify notation in this section by two conventions. First, we identify each information set with the unique node in it. Secondly, if c is a profile of strategies for the game $\Gamma(e)$ and $\Gamma(e,t)$ is a proper subgame of $\Gamma(e)$ in which the restrictions of c to $\{t\}\cup \succ_t$ are employed as strategies, then we denote player i's payoff in the subgame $\Gamma(e,t)$ by $u_i(.;\Gamma(e,t))$.

Lemma 2.1.3. *Consider a standard perfect information game in extensive form* $\Gamma(e)$*. A strategy profile* $c=(c_i)_{i\in N}$ *is not a subgame perfect equilibrium of* $\Gamma(e)$ *iff. there exist* $i\in N$, $t\in \iota^{-1}(\{i\})$ *and* $c_i'\in C_i$ *such that*

(A) $c_i'(t)\neq c_i(t)$,

(B) c_i *and* c_i' *agree on* $\iota^{-1}(\{i\})\cap \succ_t$, *and*

(C) $u_i(c;\Gamma(e,t))<u_i(c_{-i},c_i';\Gamma(e,t))$.

Proof. Clearly, if there exist $i\in N$, $t\in \iota^{-1}(\{i\})$ and $c_i'\in C_i$ such that (C) is satisfied, then c is not a subgame perfect equilibrium of $\Gamma(e)$.

Conversely, suppose c is not a subgame perfect equilibrium of $\Gamma(e)$. By definition, there exists $t^*\in T$, $i\in N$ and $c_i^*\in C_i$ such that

$$u_i(c;\Gamma(e,t^*))<u_i(c_{-i},c_i^*;\Gamma(e,t^*)) \qquad (*)$$

Let C_i^* be the set of strategies in C_i that satisfy $(*)$. For each $c_i^*\in C_i^*$,

$$\delta(c_i^*)=|\{\tau\in \iota^{-1}(\{i\})\mid c_i^*(\tau)\neq c_i(\tau) \text{ and } \tau\in\{t^*\}\cup\succ_{t^*}\}|$$

is the number of nodes in the subgame $\Gamma(e,t^*)$ at which c_i^* deviates from c_i. The finiteness of $\Gamma(e)$ and $(*)$ imply that $\delta(c_i^*)\in\mathcal{N}$ for every $c_i^*\in C_i^*$. Moreover, as the game is finite, C_i^* is finite. Consequently, there exists c_i' that minimizes δ on C_i^*.

If $\delta(c_i') = 1$, then let $t \in \iota^{-1}(\{i\})$ be the unique node such that $c_i'(t) \neq c_i(t)$ and $t \in \{t^*\} \cup \succ_{t^*}$. By construction, i, t and c_i' satisfy (A) and (B). Moreover, $(*)$ implies that player i changes the outcome in subgame $\Gamma(e, t^*)$ by using strategy c_i' instead of c_i. As c_i' deviates from c_i only at t, it follows that, if t^* is reached, c generates a path that reaches t. Consequently, (C) is satisfied as

$$u_i(c; \Gamma(e,t)) = u_i(c; \Gamma(e,t^*)) < u_i(c_{-i}, c_i'; \Gamma(e,t^*)) = u_i(c_{-i}, c_i'; \Gamma(e,t))$$

Suppose $\delta(c_i') > 1$. Let $t \in \iota^{-1}(\{i\})$ be such that $t \in \{t^*\} \cup \succ_{t^*}$ and $c_i'(t) \neq c_i(t)$, while c_i' and c_i agree on $\iota^{-1}(\{i\}) \cap \succ_t$. Such a node t exists as $\Gamma(e)$ is finite. Thus, i, t and c_i' satisfy (A) and (B). We show that $u_i(c; \Gamma(e,t)) \geq u_i(c_{-i}, c_i'; \Gamma(e,t))$ leads to a contradiction, thus establishing (C).

Suppose $u_i(c; \Gamma(e,t)) \geq u_i(c_{-i}, c_i'; \Gamma(e,t))$. Define $\hat{c}_i \in C_i$ that agrees with c_i' on $\iota^{-1}(\{i\}) - \{t\}$, while $\hat{c}_i(t) = c_i(t)$. If the play of $\Gamma(e, t^*)$ under the profile (c_{-i}, c_i') does not reach t, then the replacement of c_i' by $\hat{c}_i$ is payoff irrelevant in $\Gamma(e, t^*)$, i.e.,

$$u_i(c; \Gamma(e,t^*)) < u_i(c_{-i}, c_i'; \Gamma(e,t^*)) = u_i(c_{-i}, \hat{c}_i; \Gamma(e,t^*))$$

On the other hand, *suppose the play of* $\Gamma(e, t^*)$ *under the profile* (c_{-i}, c_i') *reaches* t. Then,

$$u_i(c; \Gamma(e,t^*)) < u_i(c_{-i}, c_i'; \Gamma(e,t^*)) = u_i(c_{-i}, c_i'; \Gamma(e,t)) \qquad (**)$$

The inequality follows from the fact that $c_i' \in C_i^*$, while the equality follows from the italicized hypothesis. Also note that

$$u_i(c; \Gamma(e,t)) = u_i(c_{-i}, \hat{c}_i; \Gamma(e,t)) = u_i(c_{-i}, \hat{c}_i; \Gamma(e,t^*)) \qquad (***)$$

The first equality follows from the fact that c and $(c_{-i}, \hat{c}_i)$ coincide on $\Gamma(e,t)$. As $(c_{-i}, \hat{c}_i)$ and $(c_{-i}, \hat{c}_i)$ differ only at t, the italicized hypothesis implies that the play of $\Gamma(e, t^*)$ under the profile $(c_{-i}, \hat{c}_i)$ reaches t. The second equality follows immediately. As $u_i(c; \Gamma(e,t)) \geq u_i(c_{-i}, c_i'; \Gamma(e,t))$, $(**)$ and $(***)$ imply that $u_i(c; \Gamma(e,t^*)) < u_i(c_{-i}, \hat{c}_i; \Gamma(e,t^*))$.

Thus, $\hat{c}_i \in C_i^*$ and $\delta(\hat{c}_i) = \delta(c_i') - 1$, a contradiction. ■

Theorem 2.1.3. *A standard perfect information game* $\Gamma(e)$ *has a subgame perfect equilibrium in pure strategies.*

Proof. For every node $t \in T$, there is a unique path from t_0 to t; let $l(t)$ denote the number of nodes in this path. Define $T_k = \{t \in T \mid l(t) = k\}$ for $k \in \mathcal{N}$. Let $n = \max_{t \in T} l(t)$, which exists as the game is finite. Clearly, $T_n \subset Z$. Define $v_n : N \times T_n \to \Re$ by $v_n(i,t) = v(i,t)$. Suppose, for $m \in \{2,\ldots,n\}$, we are given $v_m : N \times T_m \to \Re$ and $t \in T_{m-1} \cap X$. By definition, $a(t) \subset T_m$. Let

$$d_t \in \arg\max_{d \in D_t} v_m(\iota(t), \alpha_t(d))$$

and define $v_{m-1} : N \times T_{m-1} \to \Re$ by

$$v_{m-1}(i,t) = \begin{cases} v(i,t), & \text{if } t \in T_{m-1} \cap Z \\ v_m\big(i, \alpha_t(d_t)\big), & \text{if } t \in T_{m-1} \cap X \end{cases}$$

Define player i's pure strategy as follows: if $\iota(t) = i$, then $c_i(t) = d_t$.

For $t \in T_n$, $u_i(c; \Gamma(e,t)) = v(i,t) = v_n(i,t)$. Suppose that, for $t \in T_m$, $u_i(c; \Gamma(e,t)) = v_m(i,t)$. Consider $t \in T_{m-1} \cap Z$. By definition, $u_i(c; \Gamma(e,t)) = v(i,t) = v_{m-1}(i,t)$. Consider $t \in T_{m-1} \cap X$. Suppose $\iota(t) = i$ and $c'_i \in C_i$ agrees with c_i on $\iota^{-1}(\{i\}) \cap \succ_t$ and $c'_i(t) = d'_t$. Then, by the inductive hypothesis,

$$u_i(c_{-i}, c'_i; \Gamma(e,t)) = u_i(c_{-i}, c'_i; \Gamma(e, \alpha_t(d'_t))) = u_i(c; \Gamma(e, \alpha_t(d'_t))) = v_m(i, \alpha_t(d'_t))$$

Note the special case, $u_i(c; \Gamma(e,t)) = v_m(i, \alpha_t(d_t)) = v_{m-1}(i,t)$.

Consider $t \in X$. Suppose $\iota(t) = i$ and $l(t) = m$. Suppose $c'_i \in C_i$ agrees with c_i on $\iota^{-1}(\{i\}) \cap \succ_t$ and $c'_i(t) = d'_t$. Then, $u_i(c; \Gamma(e,t)) = v_{m+1}(i, \alpha_t(d_t)) \geq v_{m+1}(i, \alpha_t(d'_t)) = u_i(c_{-i}, c'_i; \Gamma(e,t))$. Lemma 2.1.2 implies that c is a subgame perfect equilibrium. ∎

Combining this result with the minmax theorem of zero-sum games yields an interesting fact about chess. Given the rule that repetition of any position thrice results in the termination of the game in a stalemate, chess is a finite game of perfect information. Theorem 2.1.3 guarantees the existence of a pure strategies Nash equilibrium for chess. Moreover, chess can be modelled as a zero-sum game if the loser gives the winner a unit of utility; both players get zero payoffs in case of a stalemate. For simplicity, we rule out the possibility of players offering draws or resigning; this is chess to the finish! Given the pure strategy equilibrium profile, the equilibrium payoff vectors must be either $(1,-1)$, or $(-1,1)$, or $(0,0)$. We know from Theorem 1.6.2(A) that the payoffs of all the players are invariant across Nash equilibria. Therefore, the payoff profile in *every* equilibrium is either (a) $(1,-1)$ (i.e., white wins in every equilibrium), or $(-1,1)$ (i.e., black wins in every equilibrium), or $(0,0)$ (i.e., every game ends in a stalemate in equilibrium). Moreover, Theorem 1.6.2(A) implies that employing the proposed equilibrium strategy guarantees the equilibrium payoffs for either player *regardless of what the other player does*! Thus, one and only one of the following statements is true: (a) white has a pure strategy that ensures a win regardless of what black does, (b) black has a pure strategy that ensures a win regardless of what white does, or (c) both players have strategies that ensure at least a stalemate regardless of what the other player does. Therefore, if we could practically implement the algorithm used in Theorem 2.1.3, then chess would be a very dull game!

2.2 Incomplete information games

We have so far studied games with the strategic form $\Gamma = \{N; (A_i, u_i)_{i \in N}\}$, implicitly assuming throughout that Γ is common knowledge among the players. We now consider more general strategic situations in which Γ is not common knowledge. The set of players N and the action sets $(A_i)_{i \in N}$ will continue to be common knowledge, but the players' utility functions $(u_i)_{i \in N}$ will not be common knowledge.[9]

This lack of common knowledge about utility functions is modelled by

• drawing each player's utility function from a parametrized family of utility functions, and

• assuming that the value of this parameter is not common knowledge.

Let T_i be player i's *type space*; $t_i \in T_i$ is player i's *type* or private information. $T = \prod_{i \in N} T_i$ is the space of *type profiles* $t = (t_i)_{i \in N}$. The parameter that will determine the players' utility functions is $t \in T$. The value of t is not common knowledge; given the type profile t, player i knows t_i but not t_j for $j \neq i$. By way of interpretation, we assume that Nature draws player i's type $t_i \in T_i$ and privately informs player i about this choice. In general, we shall assume that player i's utility depends on the types of all the players, i.e., player i's utility is $u_i : T \times A \to \Re$. We can associate with each type profile $t \in T$, the strategic form game $\Gamma(t) = \{N, (A_i, u_i(t))_{i \in N}\}$. Thus, we have a parametrized family of strategic form games $\{\Gamma(t) \mid t \in T\}$, where the parameter t is not common knowledge; this family will be referred to as an *incomplete information game*. In order to study such a game, we need a way of modelling each player's decisions as the solution of some well-

[9] By appropriate transformations, it is possible to convert formally the lack of common knowledge about the other aspects of the game to the assumed lack of common knowledge about utility functions.

formulated decision problem. We proceed to develop the Bayesian model of incomplete information games.

2.3 The Bayesian model

Formulating the Bayesian model involves supplementing the given incomplete information game with structure that enables a sensible modelling of the informational and incentive aspects of the game.

Definition 2.3.1. *An abstract Bayesian game is a collection*

$$\Gamma(B) = \{N, ((\Omega, \mathcal{F}, \mu_i), (T_i, \mathcal{T}_i), \Theta_i, (A_i, \mathcal{A}_i), u_i)_{i \in N}\}$$

N is the set of players and $(\Omega, \mathcal{F})$ is a measurable state space. For each player $i \in N$, $\mu_i \in \mathcal{P}(\Omega)$ is his belief about the state of the world; $(T_i, \mathcal{T}_i)$ is his measurable type space; $\theta_i : \Omega \to T_i$ is a surjective $\mathcal{F}/\mathcal{T}_i$ measurable random variable; $(A_i, \mathcal{A}_i)$ is his measurable action space; and $u_i : T \times A \to \Re$ is his utility function.

Since players do not know each others types, it is natural to model the type profile $t \in T$ as the outcome of a random variable defined on some state space; given the state $\omega \in \Omega$, player i's type is $\theta_i(\omega)$, which is privately communicated by Nature to player i. Given this interpretation, the σ-algebra $\mathcal{T}_i$ is a model of player i's ability to distinguish between various type-signals received from Nature. Assuming surjectivity of θ_i involves no loss of generality and will be useful when we construct posterior beliefs for the players conditional on their private information. Define $\theta : \Omega \to T$ by $\theta(\omega) = (\theta_i(\omega))_{i \in N}$; note the relations $\theta_i = \pi_i \circ \theta$ and $\theta_{-i} = \pi_{-i} \circ \theta$.

Since player i's information about the state is derived from the private signal θ_i, his information about the state is represented by the σ-algebra $\sigma(\theta_i) = \{\theta_i^{-1}(E) \mid E \in \mathcal{T}_i\}$. Clearly, $\sigma(\theta_i)$ is the smallest σ-algebra on Ω that makes the function θ_i a random variable.

Given an incomplete information game, player i's action can be made contingent on his own type; thus, a player's strategy is a $\mathcal{T}_i/\mathcal{A}_i$ measurable function $s_i : T_i \to A_i$; let S_i be the set of all such functions. Define $S = \prod_{i \in N} S_i$. Player i's utility depends on Nature's choice of t and the players' choice of a. When Nature chooses state ω, the players' types are $t = \theta(\omega)$, while the resulting action profile is $a = s(t) = s \circ \theta(\omega)$. Therefore, in order to choose a strategy, not only must a player have a belief about the state ω, but must also have a belief about the mapping θ that generates the type profile and the strategies s_{-i} that generate the actions of the other players. This problem is simplified by the following axiom.

Assumption 2.3.2. *The mappings $\theta : \Omega \rightarrow T$ and $s : T \rightarrow A$ are common knowledge.*

This assumption is not merely a formal simplification. We wish to interpret player i's type t_i as a *comprehensive* description of data that is known to player i but is not common knowledge; if the strategies were not common knowledge, then t_i would be an inadequate description of such data.

Assumption 2.3.2 implies that the only uncertainty facing a player is the identity of the state of the world. Given the abstract Bayesian game $\Gamma(B)$, define $U_i : S \rightarrow \Re$ by

$$U_i(s) = \int_\Omega \mu_i(d\omega) u_i(\theta(\omega), s \circ \theta(\omega))$$

Player i's *ex ante* decision problem in $\Gamma(B)$ is to choose $s_i \in S_i$ to maximize $U_i(s_{-i}, s_i)$.

Definition 2.3.3. *Given the abstract Bayesian model $\Gamma(B)$, a Bayesian equilibrium is a profile of functions $s = (s_i)_{i \in N} \in S$ such that, for every $i \in N$ and $s'_i \in S_i$,*

$$U_i(s) \geq U_i(s_{-i}, s'_i)$$

Clearly, the abstract Bayesian game $\Gamma(B)$ generates an imperfect information game in which

(1) Nature randomly selects a type profile t and privately informs each player i about his selected type t_i, and

(2) all players simultaneously choose their actions based on this private information.

Given the definitions of $(S_i, U_i)_{i \in N}$, it is clear that a Bayesian equilibrium of $\Gamma(B)$ is nothing but a Nash equilibrium of the strategic form game $\Gamma^* = \{N, (S_i, U_i)_{i \in N}\}$.

2.4 The canonical Bayesian model

Consider the abstract Bayesian game $\Gamma(B)$ of Definition 2.3.1. Given player i's belief $\mu_i \in \mathcal{P}(\Omega)$, θ has the distribution $P_i = \mu_i \circ \theta^{-1} \in \mathcal{P}(T)$. If players choose the strategies $s = (s_j)_{j \in N}$, then player i's expected utility can be written as

$$\int_\Omega \mu_i(d\omega) u_i(\theta(\omega), s \circ \theta(\omega)) = \int_T \mu_i \circ \theta^{-1}(dt) u_i(t, s(t)) = \int_T P_i(dt) u_i(t, s(t))$$

Consequently, the abstract probability space $(\Omega, \mathcal{F}, \mu_i)$ can be replaced by the canonical probability space $(T, \mathcal{T}, P_i)$, and the abstract Bayesian game $\Gamma(B)$ can be replaced by its canonical version.

Definition 2.4.1. *A canonical Bayesian game is a collection*

$$\Gamma(B) = \{N, ((T, \mathcal{T}, P_i), (T_i, \mathcal{T}_i), \pi_i, A_i, u_i)_{i \in N}\}$$

Note that the projection $\pi_i : T \to T_i$ generates player i's private information given the canonical state $t \in T$. We shall henceforth adhere to the canonical version of the Bayesian model.

2.5 Conceptual issues

The heterogeneity of beliefs allowed in Definitions 2.3.1, 2.3.3 and 2.4.1 is controversial. We turn to this issue now.

The so-called Harsanyi doctrine insists that differences in beliefs across players must be attributable to differences in information. More formally, if beliefs are different across players, these beliefs must be posterior beliefs derived from some common prior and conditional on different information. If one subscribes to this doctrine, apparent *a priori* differences in beliefs merely mean that the notions of "state of the world" and "information" are misspecified. While it is possible to defend this reductionist position as a truism with sufficiently general definitions of "state of the world" and "information", the more important defense of this doctrine is methodological. With *a priori* differences in beliefs *and* information, it is impossible to systematically disentangle their effects. If differences in beliefs arise endogenously within a model because of informational asymmetries, the results of the model can be attributed entirely to the informational asymmetries. The implication of accepting the Harsanyi doctrine is the following assumption, which we shall maintain throughout.

Assumption 2.5.1. *(Common priors axiom) The players have a common prior belief* $P \in \mathcal{P}(T)$.

Even if we subscribe to the Harsanyi doctrine, actual application of this doctrine to concrete strategic situations can be problematic. In many situations, it is reasonable to assume that there is a common belief for all sensible players to hold; e.g., where the probability distribution of the state is commonly known, as in the case of a known fair coin being tossed. In many situations, however, the state may be generated by an experiment whose probability law is not commonly known, e.g., in a situation in which bets are placed on the outcome of a horse race. One can argue in line with the Harsanyi doctrine that the players hold different opinions because they must have different information about the horses, the weather conditions; indeed, the genetic information of entirely sensible players might affect their beliefs about the race. Arguing that the notions of the state and information are misspecified in such cases does not clinch the argument for the Harsanyi doctrine. The clinching argument needs to show that there is indeed a probability space and a notion of information that will rationalize the differences in beliefs. Mertens and Zamir (1985) provided this argument by constructing the canonical probability space that generates the heterogeneous beliefs. Thus, in principle, all situations with heterogeneous beliefs can be transformed into the correct homogeneous belief situation by using the Mertens-Zamir technique.

Very roughly, the Mertens-Zamir argument runs as follows. Suppose players have different beliefs defined an a given 'state' space T. Since these beliefs must have arisen because of different private information, the beliefs themselves must be private information. Thus, each player's type must also include a description of his belief on T. This amounts to extending the state space to $T^1 = T \times \mathcal{P}(T)^N$, so that in the extended state $t^1 = (t, (P_i)_{i \in N})$, player i's type is $t_i^1 = (t_i, P_i)$. The problem does not end here as players could have subjective beliefs on T^1 too. As should be obvious, this extension of the state space and the notion of types can go on indefinitely, unless we reach a finite stage where we suddenly have a common prior. Since such a stage may not exist, one must be able to deal with an infinite hierarchy of beliefs over a nested hierarchy of type spaces. Is there a 'universal' probability space and a notion of information such that the entire hierarchy of beliefs can be rationalized as conditional distributions given the prior belief and some information process? Mertens and Zamir showed that it is possible to construct such a universal probability space.

9.5 Existence of Bayesian equilibria

The next important issue is whether Bayesian equilibria exist in a wide enough class of Bayesian games. The following theorem identify a fairly rich class of games that possess Bayesian equilibria.

Theorem 9.5.1. *Consider the canonical Bayesian game $\Gamma(B)$ given in Definition 2.4.1. If*

(a) N is finite, and

for every $i \in N$,

(b) $A_i \subset \Re^l$ is nonempty, compact and convex, for some $l < \infty$,

(c) T_i is finite,

(d) $u_i(t, .) : A \to \Re$ is continuous, and

(e) $u_i(t, a_{-i}, .) : A_i \to \Re$ is concave for every $t \in T$ and $a_{-i} \in A_{-i}$,

then $\Gamma(B)$ has a Bayesian equilibrium.

Proof. Since T_i is finite, we shall endow it with the discrete σ-algebra and the discrete topology; consequently, every subset of T_i is measurable and open. It follows that every function $s_i : T_i \to A_i$ is measurable and continuous. Player i's strategy space is $S_i = A_i^{T_i}$; all functions in S_i are measurable and continuous. Clearly, S_i is nonempty. By Tychonoff's theorem, S_i is compact. Elementary arguments show that S_i is a convex subset of $\Re^{l \times |T_i|}$.

Player i's expected utility is $U_i : S \to \Re$, given by

$$U_i(s) = \int_T P_i(dt) u_i(t, s(t)) = \sum_{t \in T} P_i(t) u_i(t, s(t))$$

It follows from the concavity of $u_i(t, a_{-i}, .)$ that $U_i(s_{-i}, .) : S_i \to \Re$ is concave, and therefore quasi-concave. To check continuity, let $(s^n) \subset S$ be a sequence that converges to $s \in S$. It follows that, for every $t \in T$, $(s^n(t))$ converges to $s(t)$, and therefore by the continuity of $u_i(t, .)$, $(u_i(t, s^n(t)))$ converges to $u_i(t, s(t))$. It follows that $(U_i(s^n))$ converges to $U_i(s)$. Now the result follows from Nash's existence theorem. ■

3. Application I: Auction design

3.1 Introduction

Auctions are perhaps the most familiar allocation mechanism we know. They are used not only to sell but also to buy, e.g., procurement by governments and firms *via* a tendering process. We shall restrict attention to the process of selling by auction; there is an analogous theory of buying via auctions.

The theory of auction design for a single object is extremely rich, while the theory of auction design for a sequence of objects is less settled and still quite open. We shall concentrate on the classical theory related to selling a single object. This theory has two aspects. First, there is the positive theory, which characterizes the equilibria that can arise for different auction mechanisms. Secondly, there is the normative theory, which characterizes the nature of the optimal auction mechanism. As this theory is quite vast, we shall illustrate these two themes by studying them in the context of the simplest possible auction environment. Most of the results that we derive in the context of our simple example can be generalized, often with little effort, and sometimes a little more.

3.2 Sealed bid second price (Vickrey) auction

Consider a seller with a single indivisible object facing two potential buyers of that object, say Buyer 1 and Buyer 2. Each buyer knows his own valuation of the object but not that of the other buyer. The seller does not know the valuation of either buyer. However, it is common knowledge that each buyer's valuation is either 3 or 4 with equal probability. Moreover, the valuations of the two buyers are independent, i.e., knowing the valuation of one buyer yields no information about the valuation of the other buyer. This is the environment we shall carry with us throughout the discussion of auctions.

Suppose the seller runs a sealed bid second price auction (i.e., a Vickrey auction). Both buyers simultaneously submit sealed bids to the seller. The seller awards the object to the highest bidder who pays the amount bid by the second highest bidder. If they bid the same amount, then the seller awards the object with equal probability to either buyer, who pays the amount of the common bid. A buyer gets payoff 0 if he does not get the object, and gets payoff equal to the difference between his valuation and his payment if he wins the object.

Note that this procedure generates an incomplete information game, with each buyer's valuation being his private information. Consequently, a buyer's strategy in the game so generated has to specify the bid made by that buyer conditional on his valuation.

Theorem 3.2.1. *It is a weakly dominant strategy for each buyer to submit a bid equal to his valuation, i.e., the Vickrey auction induces truth-telling as a dominant strategy.*

Proof. Consider Buyer 1; the argument for Buyer 2 is analogous. His valuation is either 3 or 4. We show that, in either case, bidding the true value does at least as well as any other bid, against any bid made by the other buyer, i.e., bidding the true value is a weakly dominant strategy. Note that, bidding the true value guarantees that the payoff is not negative; thus, if an alternative bid yields a negative payoff, it can be strictly improved upon by bidding the true value.

Suppose Buyer 1's valuation is 3. Let Buyer 1's bid be b. Let c be Buyer 2's bid. There are three possible cases: $b > c$, $b < c$ and $b = c$.

If $b > c$, then Buyer 1 wins the object and pays c for it; thus, Buyer 1's payoff is $3 - c$. If $3 \geq c$, then bidding the true valuation 3 yields exactly the same payoff as the bid b. If $3 < c$, then the payoff is negative and the above Observation applies.

If $b < c$, then Buyer 1 loses the object and gets payoff 0. If $3 > c$, then bidding the true valuation 3 would have enabled Buyer 1 to win the object and get payoff $3 - c > 0$, thus doing strictly better than by bidding b. If $3 \leq c$, then bidding 3 yields payoff 0, which is the same as from bidding b.

If $b = c$, then Buyer 1 wins the object with probability 0.5. Thus, his expected payoff is $0.5 \times (3 - b) + 0.5 \times 0 = 0.5 \times (3 - b) = 0.5 \times (3 - c)$. If $b = c > 3$, then Buyer 1's payoff is negative and the above Observation applies. If $b = c < 3$, then bidding the true valuation 3 would have enabled Buyer 1 to win the object with probability 1, thus getting payoff $3 - c$ rather than $0.5 \times (3 - c)$.

An analogous argument holds when Buyer 1's valuation is 4. ■

Theorem 3.2.2. *The expected payoff of either buyer, given that his valuation is 3, is 0. The expected payoff of either buyer, given that his valuation is 4, is* 0.5. *The expected payoff of either buyer, before knowing his valuation, is* 0.25.

Proof. Consider Buyer 1. Suppose his valuation is 3. Given the tie breaking rule and the weakly dominant bidding strategies, the probability of Buyer 1 winning is 0.25, and the payoff from winning is 0. So, the expected payoff is 0.

Suppose his valuation is 4. Given the dominant strategies, the only time he gets a positive payoff is when Buyer 2's valuation is 3. In this case, Buyer 1 gets payoff 1. Since the probability of this event is 0.5, Buyer 1's expected payoff is 0.5.

It follows from the above two arguments that, before knowing his valuation, Buyer 1's expected payoff is $0.5 \times 0 + 0.5 \times 0.5 = 0.25$. ■

Theorem 3.2.3. *The seller's expected revenue from the Vickrey auction is* 3.25.

Proof. The seller's revenue is 4 when both buyers have valuation 4, and 3 otherwise. Therefore, his expected revenue is $0.25 \times 4 + 0.75 \times 3 = 3.25$. ■

3.3 Sealed bid first price auction

Suppose the seller runs a sealed bid first price auction. Both buyers simultaneously submit sealed bids to the seller. The seller awards the object to the buyer with the higher bid, who pays the amount of his bid. If they bid the same amount, then the seller awards the object with equal probability to either buyer, who pays the amount of the common bid. The payoffs are as specified in Section 3.2. As in Section 3.2, a buyer's strategy in the game generated by the sealed bid first price auction mechanism has to specify the bid made by that buyer conditional on his valuation.

Theorem 3.3.1. *There is no Nash equilibrium in pure strategies.*

Proof. Suppose there is a pair of pure strategies $(b(.), c(.))$ that constitute a Nash equilibrium; say, Buyer 1's strategy is given by $b(3) = b_3$ and $b(4) = b_4$, and Buyer 2's strategy is given by $c(3) = c_3$ and $c(4) = c_4$. We show that this leads to a contradiction, and therefore, no such equilibrium exists.

We first show that $b_3 \leq 3$ *and* $c_3 \leq 3$. Suppose $b_3 > 0$. It follows that $c_3 > b_3$ and $c_4 > b_3$; otherwise, Buyer 1 will win with a positive probability and get a negative payoff, which can be unilaterally improved upon by setting $b(3) = 3$. But this means, Buyer 2 wins by bidding $c_3 > 3$ when his valuation is 3, leading to a negative payoff; this can be improved upon by setting $c(3) = 3$. So, if $(b(.), c(.))$ is a Nash equilibrium, then $b_3 \leq 3$ and $c_3 \leq 3$.

By an analogous argument, $b_4 \leq 4$ *and* $c_4 \leq 4$.

Next, we show that $b_4 > 3$ *and* $c_4 > 3$. Suppose $b_4 \leq 3$. It follows that $c_4 < b_4 \leq 3$; if $b_4 \leq c_4 < 4$, then Buyer 1 has an incentive to unilaterally deviate by setting $b(4) \in (c_4, 4)$; if $b_4 < c_4 = 4$, then Buyer 1 has an incentive to unilaterally deviate by setting $b(4) = 4$. But, if $c_4 < b_4 \leq 3$, then Buyer 2 would like to unilaterally deviate to $c(4) > 3$. Thus, if $(b(.), c(.))$ is a Nash equilibrium, then $b_4 > 3$ and $c_4 > 3$.

Next, we show that $b_3 = 3$ *and* $c_3 = 3$. If $b_3 < c_3 \leq 3$, then Buyer 1 (and Buyer 2) have obvious profitable deviations; note that this argument can be made unambiguously because $b_4 > 3$ and $c_4 > 3$. If $b_3 = c_3 < 3$, then both buyers have obvious profitable deviations. So, we must have $b_3 = 3$ and $c_3 = 3$.

Next, we show that $b_4 = 4$ *and* $c_4 = 4$. If $3 < b_4 < c_4 \leq 4$, then either buyer has a profitable unilateral deviation. So, we must have $3 < b_4 = c_4 \leq 4$. Suppose $3 < b_4 = c_4 < 4$. Again, either buyer has a profitable unilateral deviation. The only possibility left is $b_4 = c_4 = 4$.

Finally, $b_3 = c_3 = 3$ *and* $b_4 = c_4 = 4$ *does not yield a Nash equilibrium.* This is because either buyer has an incentive to unilaterally deviate to a lower bid when his valuation is 4. ■

Theorem 3.3.2. *Neither buyer has a weakly dominant pure strategy. In particular, truth-telling is not a weakly dominant strategy.*

Proof. If either buyer has a weakly dominant pure strategy, so does the other. This pair of identical weakly dominant strategies would be a Nash equilibrium, which is ruled out by Claim 1. ■

Theorem 3.3.3. *There is a Nash equilibrium in mixed strategies.*

Proof. Suppose both buyers use the following strategy. If the valuation is 3, bid 3. If the valuation is 4, choose a bid between 3 and 3.5 using a randomization device such that the probability of a bid less than or equal to x is $(x - 3)/(4 - x)$.

Note: (a) the probability of a bid less than 3 or greater than 3.5 is 0, and (b) the randomization strategy amounts to specifying the distribution function

$$F(x) = \begin{cases} 0, & \text{if } x < 3 \\ (x-3)/(4-x), & \text{if } 3 \leq x < 3.5 \\ 1, & \text{if } x \geq 3.5 \end{cases}$$

We show that the proposed pair of strategies is a Nash equilibrium.

Consider Buyer 1; an analogous argument applies to Buyer 2. Suppose Buyer 1's valuation is 3. We show that, given Buyer 2's strategy, Buyer 1 cannot do better than by bidding 3. Note that, bidding the true valuation guarantees at least 0 payoff.

With probability 0.5, Buyer 2's valuation is 3; in this case, Buyer 2's strategy requires him to bid 3. With probability 0.5, Buyer 2's valuation is 4; in this case, Buyer 2's strategy requires him to randomize over the interval $[3, 3.5]$ with distribution F.

If Buyer 1 bids $b < 3$, then he has no chance of winning the object. Therefore, his expected payoff is 0, which is no better than the guaranteed non-negative payoff from bidding the true valuation 3.

If Buyer 1 bids $b > 3$, then he has a probability greater than 0.5 of winning the object and getting the payoff $3 - b < 0$. Since such a bid implies a negative expected payoff, it is better to bid 3, which guarantees a non-negative payoff.

Suppose Buyer 1's valuation is 4 and he bids b with $3 \leq b \leq 3.5$. If he wins the object, his payoff is $4 - b$, and if he loses the auction his payoff is 0. Thus, his expected payoff is $4 - b$ times the probability of winning the object. We calculate this probability. With probability 0.5, Buyer 2's valuation is 3, and his strategy will require him to bid 3, which implies Buyer 1 will win the object. With probability 0.5, Buyer 2's valuation is 4, and his strategy will require him to randomize his bids such that the probability of bidding less than or equal to b is $(b-3)/(4-b)$. Thus, the probability of Buyer 1 winning the object is

$$0.5 + 0.5 \times \frac{b-3}{4-b} = 0.5 \times \frac{1}{4-b}.$$

It follows that Buyer 1's expected payoff from bidding b such that $3 \leq b \leq 3.5$, when his own valuation is 4, is

$$0.5 \times \frac{1}{4-b} \times (4-b) = 0.5.$$

Since this expected payoff is independent of b when $3 \leq b \leq 3.5$, the expected payoff from any randomization over this interval will also yield exactly the same expected payoff.

By bidding higher than 3.5, i.e., by choosing $b > 3.5$, Buyer 1 can *ensure* that he wins the object when his valuation is 4, but this is a Pyrhic victory as the resulting payoff $4 - b$ will be less than 0.5, which he gets from the strategy described above!

By bidding less than 3, Buyer 1's probability of winning the object becomes 0, and consequently, his expected payoff is 0. Thus, the prescribed strategy for Buyer 1 is a best response to the (identical) strategy prescribed for Buyer 2. Since an analogous argument applies to Buyer 2, the prescribed pair of strategies is a Nash equilibrium. ■

Theorem 3.3.4. *Given the Nash equilibrium described above, the expected payoff of either buyer, given that his valuation is 3, is 0; the expected payoff of either buyer, given that his valuation is 4, is* 0.5*; the expected payoff of either buyer before knowing his valuation is* 0.25*.*

Proof. Consider Buyer 1. Suppose his valuation is 3. The equilibrium strategy requires him to bid 3. Since the resulting payoff is 0 in all circumstances, the expected payoff is 0.

Suppose his valuation is 4. The calculation of Buyer 1's expected payoff, given the strategy of Buyer 2, is contained in the proof of Claim 3.

It follows from the above two arguments that, *before* knowing his valuation, Buyer 1's expected payoff is $0.5 \times 0 + 0.5 \times 0.5 = 0.25$. ■

Theorem 3.3.5. *The expected revenue of the seller is* 3.25*.*

Proof. Suppose both buyers have valuation 3. Then, they will both bid 3, resulting in a revenue equal to 3. Thus, the seller's expected revenue, conditional on both buyers having valuation 3, is 3.

Suppose one buyer has valuation 3 and the other has valuation 4. Since the buyer with valuation 3 will bid 3, the winner will be the buyer with valuation 4, and the seller's revenue will equal this buyer's bid. Since the buyer with valuation 4 will be randomizing, the seller's expected revenue will equal this buyer's expected bid. Thus, the seller's expected revenue, conditional on one buyer having valuation 3 and the other having valuation 4, is

$$\int_3^{3.5} x\, dF(x) = \int_3^{3.5} dx\, x/(4-x)^2.$$

Suppose both buyers have valuation 4. Therefore, they will be randomizing using the distribution F. We first calculate the expected revenue, given that Buyer 1 wins the auction. Suppose Buyer 1 wins the auction by bidding x. Clearly, Buyer 2 must have bid less than x. The probability of Buyer 2 bidding less than x is $(x-3)/(4-x)$. Thus, the seller's expected revenue, conditional on both buyers having valuation 4 and Buyer 1 winning the object by bidding x, is $x(x-3)/(4-x)$. Consequently, the seller's expected revenue, conditional on both buyers having valuation 4 and Buyer 1 winning the object, is

$$\int_3^{3.5} \frac{x(x-3)}{4-x}\, dF(x) = \int_3^{3.5} dx\, \frac{x(x-3)}{(4-x)^3}$$

The seller's expected revenue, conditional on both buyers having valuation 4 and Buyer 2 winning the object, is exactly the same. Thus, the seller's expected revenue, conditional on both buyers having valuation 4 and either Buyer 1 or Buyer 2 winning the object, is

$$2 \times \int_{3}^{3.5} dx \, \frac{x(x-3)}{(4-x)^3}$$

Thus, collecting together the above arguments, the seller's expected revenue is

$$0.25 \times 3 + 0.5 \times \int_{3}^{3.5} dx \, \frac{x}{(4-x)^2} + 0.25 \times 2 \times \int_{3}^{3.5} dx \, \frac{x(x-3)}{(4-x)^3}$$

It is now a straightforward, though tedious, calculation to confirm that this expected revenue is 3.25. ■

3.4 Ascending (English) auction

This requires the buyers to compete for the object by announcing sequentially higher bids. A higher bid by a competing buyer supersedes all previous lower bids. The bidding ends when no one wishes to raise the bid above the last announced bid. The object is awarded to the person who makes this last (highest) bid and he pays the amount of the announced bid.

While the details of the bidding procedure in an English auction can be complicated and vary substantially, they are analytically unimportant. We adopt the following simple procedure that captures the spirit of the English auction. The seller *continuously* raises the price. As the price rises, a buyer can opt out of the auction. The auction ends when only one buyer is left in the auction; the object is awarded to this buyer, who pays the price at which the auction stopped. If all the remaining buyers in the auction opt out together, then the object is awarded with equal probability to each of these remaining buyers; the price paid is the bid at which the auction stopped.

Given the above-described procedure, a strategy for a buyer amounts to selecting a trigger price, i.e., the maximum price the buyer is willing to pay for the object: the buyer will stay in the auction as long as the bid is below the trigger price and will opt out when the trigger price is reached. The strategic tension is obvious: opting out 'too soon' could mean having lost an opportunity to get the object at a favorable price, while staying in 'too long' could result in the winner's curse of winning the object at too high a price.

Clearly, it does not make sense to have a trigger price below one's valuation, thereby foregoing some chance of getting the object at a price less than one's valuation. Similarly, it does not make sense to have a trigger price that is higher than one's valuation as that involves some risk of having to buy the object at a price higher than one's valuation. Therefore, the right trigger price is one's valuation, i.e., it is a weakly dominant strategy for both buyers to set their trigger price equal to their valuations.

Consequently, the price of the object and the seller's revenue will coincide with the second highest valuation.

Note that the selection of a trigger price amounts to selecting a sealed bid representing the highest price at which one is willing to buy the object. The object will be won by the buyer with the highest trigger price, i.e., the highest sealed bid. However, the price paid by this buyer will be the second highest valuation. Thus, the outcome is isomorphic to a Vickrey auction. In our example with two bidders whose valuations can be 3 or 4 with equal and independent probabilities, the seller's expected revenue is 3.25.

3.5 Descending (Dutch) auction

In this auction procedure, the seller starts by announcing a 'high' selling price (for our purposes, this will be taken to be higher than the highest possible valuation across the buyers) and then continuously drops this price until some buyer accepts the last announced price. The object is awarded to that buyer at the last announced price.

A strategy for a buyer amounts to selecting a trigger price: a price he will accept. The strategic tension is obvious: hit the price 'too soon' means paying a 'high' price, while waiting for the price to drop further involves the risk of someone else hitting the price first.

Note that the selection of a trigger price amounts to selecting a sealed bid at which one is willing to buy the object. It follows that the object will be won by the buyer with the highest trigger price, i.e., the highest sealed bid, who will pay the amount of the trigger price. Thus, the situation is isomorphic to a first price sealed bid auction.

Consequently, there is no dominant strategy for any player, nor can there be a pure strategies Nash equilibrium. There is a mixed strategies Nash equilibrium that is identical to the equilibrium for the sealed bid first price auction. Consequently, in our example with two bidders whose valuations can be 3 or 4 with equal and independent probabilities, the seller's expected revenue is 3.25.

3.6 Conclusions of positive theory

• Auction mechanism significantly affects bidding strategies.

• First price auction and Dutch auction do not induce truth telling. Contrary to intuition, the bidding in a Nash equilibrium is far below valuations.

• Expected revenue is the same in all the considered cases, although the Vickrey and English auction yield more compelling predictions as buyers have dominant strategies.

• Is there a systematic way of comparing the expected revenues from different auction mechanisms in any given situation?

• More ambitious: is it possible to design an optimal auction, i.e., one that maximizes the seller's expected revenue?

• Surprisingly, the answer to both questions is yes! The key tool in both cases is the 'revelation principle'.

3.7 Optimal auction: preliminaries

• When we discuss 'optimality' of anything, one must have in mind the class of objects over which one is optimizing. So, what is the class of auctions over which one is optimizing? How is this class represented?

• If the seller is allocating a single indivisible object, one needs to specify the complete class of methods by which the seller can allocate this object. There are two elements to the description of such methods: (1) the set of messages the buyers can send to the seller (i.e., the *'language'* employed by the method), and (2) the rule by which the seller allocates the object and extracts payments, based on the messages received from the buyers (the *'outcome function'*). Since the class of allocation methods, determined by the languages and outcome functions available to the seller, is extremely vast, how does one begin to search for an optimal allocation method?

• **Crude organizing principle: Consequentialism.** We judge an allocation mechanism by the consequences (or outcomes) that result from its employment.

• What are the consequences of employing a given allocation mechanism? The consequences are the 'equilibria' of the game induced by the mechanism.

• **Refined organizing principle: Equilibrium consequentialism.** We judge an allocation mechanism by the **equilibrium** consequences that result from its employment.

• What does "consequence" mean? In the auction context, the consequence (or outcome) is the (possibly random) allocation of the object among the buyers and the (possibly random) extraction of revenue from the buyers.

• **Implication of equilibrium consequentialism.** Two allocation mechanisms that yield identical equilibrium outcomes for buyers of all possible valuations must be deemed to be equivalent.

• **Second organizing principle: Revelation principle.** Given any allocation mechanism and any equilibrium associated with it, it is possible to: (a) replace it with an equivalent allocation mechanism that is 'simple' (or 'direct'), and (b) that induces truthful revelation of valuations in equilibrium.

• **Implication of Revelation principle.** We can restrict the search for an optimal allocation mechanism to direct mechanisms that induce truth-telling in equilibrium. This yields a well-defined class of mechanisms and drastically simplifies the process of searching for an optimal auction.

3.8 Consequentialism and revelation principle at work

There are two buyers, 1 and 2, whose valuations θ_1 and θ_2 respectively, are private information. Let θ_1 and θ_2 be independent random variables that take values 3 and 4 with equal probabilities.

In a direct mechanism, the message spaces of both buyers is $\{3, 4\}$. Let $x_i(t, t')$ be the probability of Buyer i getting the object if Buyer 1 announces valuation t and Buyer 2 announces valuation t'. Clearly, $x_i(t, t') \geq 0$ for $i = 1$, 2, and $x_1(t, t') + x_2(t, t') \leq 1$; this leaves open the possibility that the seller might decide not to sell the good on receipt of certain messages. Let $T_i(t, t')$ be the transfer from Buyer i to the seller if Buyer 1 announces valuation t and Buyer 2 announces valuation t'. We assume that x_1 and x_2, and T_1 and T_2, are symmetric across the players, i.e., $x_1(t, t') = x_2(t', t)$ and $T_1(t, t') = T_2(t', t)$ for all t, $t' \in \{3, 4\}$.

We focus our attention on Buyer 1; the argument for Buyer 2 is analogous. Buyer 1's payoff if message t is sent by Buyer 1 and message t' is sent by Buyer 2 is

$$u_i(t, t', \theta_i) = \theta_i x_i(t, t') - T_i(t, t').$$

Let

$$\underline{X} = 0.5x_1(3,3) + 0.5x_1(3,4)$$
$$\bar{X} = 0.5x_1(4,3) + 0.5x_1(4,4)$$
$$\underline{T} = 0.5T_1(3,3) + 0.5T_1(3,4)$$
$$\bar{T} = 0.5T_1(4,3) + 0.5T_1(4,4).$$

$\underline{X}$ (resp. $\bar{X}$) is the probability of Buyer 1 getting the object if his valuation is 3 (resp. 4) and the buyers announce valuations truthfully. $\underline{T}$ (resp. $\bar{T}$) is the expected payment that Buyer 1 makes to the seller conditional on his valuation being 3 (resp. 4). Note that, if Buyer 1's valuation is 3, then his expected payoff conditional on this information is

$$\begin{aligned} 0.5u_1(3,3,3) + 0.5u_1(3,4,3) &= 0.5[3x_1(3,3) - T_1(3,3)] + 0.5[3x_1(3,4) - T_1(3,4)] \\ &= 3\underline{X} - \underline{T}. \end{aligned}$$

Similarly, if Buyer 1's valuation is 4, then his expected payoff conditional on this information is

$$\begin{aligned} 0.5u_1(4,3,4) + 0.5u_1(4,4,4) &= 0.5[4x_1(4,3) - T_1(4,3)] + 0.5[4x_1(4,4) - T_1(4,4)] \\ &= 4\bar{X} - \bar{T}. \end{aligned}$$

Therefore, if Buyer 1's valuation is 3, he will participate in the mechanism if and only if

$$3\underline{X} - \underline{T} \geq 0. \tag{A}$$

Similarly, if Buyer 1's valuation is 4, he will participate in the mechanism if and only if

$$4\bar{X} - \bar{T} \geq 0. \tag{B}$$

If Buyer 1's valuation is 3, truth-telling is incentive-compatible if and only if

$$3\underline{X} - \underline{T} \geq 3\bar{X} - \bar{T}. \tag{C}$$

Truth-telling is incentive-compatible for a buyer with valuation 4 if and only if

$$4\bar{X} - \bar{T} \geq 4\underline{X} - \underline{T}. \tag{D}$$

The sellers objective is to maximize $0.5\underline{T} + 0.5\bar{T} = 0.5(\underline{T} + \bar{T})$ by selecting $\underline{X}$, $\bar{X}$, $\underline{T}$, and $\bar{T}$, subject to the above constraints and the further constraints:

$$\underline{X},\ \bar{X} \in [0,1], \tag{E}$$

i.e., the probabilities of getting the object conditional on valuations lie between 0 and 1;

$$\underline{T},\ \bar{T} \geq 0 \tag{F}$$

i.e., the conditional payments to the seller are non-negative;

$$0.5\underline{X} + 0.5\bar{X} \leq 0.5 \tag{G}$$

or equivalently, $\underline{X} + \bar{X} \leq 1$, i.e., the prior probability of getting the object is not greater than 0.5 (it can be less than 0.5 because the seller may decide to not sell it);

$$\underline{X} \leq 0.5 + 0.5/2 = 0.75 \qquad \text{and} \qquad \bar{X} \leq 0.5 + 0.5/2 = 0.75 \tag{H}$$

i.e., a buyer's probability of getting the object cannot exceed the probability of the other buyer being of the other type plus half the probability that the other buyer is of the same type. Note that both the constraints in (H) cannot be binding because this violates (G).

We start by analyzing the constraints. Consider a solution such that $\underline{X}$, $\bar{X} > 0$. Then, (A) and (D) are binding, and (B) is not binding. Using (A) and (D), the objective function can be written as $(3 - 0.5 \times 4)\underline{X} + 0.5 \times 4\bar{X} = \underline{X} + 2\bar{X}$. Since the optimum involves $\underline{X} > 0$, we must have $3 > 0.5 \times 4$; otherwise, it is optimal to set $\underline{X} = 0$. Furthermore, at most one of the two constraints in (H) can be binding; otherwise, (G) is violated. Since we are considering an optimum, exactly one of the two constraints in (H) is binding. Thus (G) must be binding at the optimum, i.e., $0.5\underline{X} + 0.5\bar{X} = 1/2$ or $\underline{X} + \bar{X} = 1$. Using this, the objective function of the principal is

$$(3 - 0.5 \times 4) + (4 - 3)\bar{X} = 1 + \bar{X}.$$

Clearly, the optimum choice is $0.5 + 0.5/2 = 0.75$. Using the constraint (G), we have $\underline{X} = 0.25$.

Since (A) is binding, we have $\underline{T} = 3\underline{X} = 0.75$. Since (D) is binding, we have $\bar{T} = \underline{T} + 4(\bar{X} - \underline{X}) = 0.75 + 4 \times 0.5 = 2.75$. Thus, the seller's expected revenue from Buyer

1 is $0.5(\underline{T} + \bar{T}) = 1.75$. By symmetry, the seller's expected revenue from Buyer 2 is 1.75. Thus, the seller's total expected revenue from both buyers is 3.5. This is the maximum expected revenue the seller can extract from the buyers.

So, is there an auction procedure that yields an expected revenue of 3.5 in equilibrium? Yes!

3.9 An optimal auction procedure

• **Modified Vickrey auction.** The buyers are asked to submit bids from the set $\{3, 4\}$. If the bids are the same, the object is awarded with equal probability to either buyer and the winner pays the seller the amount of the bid. If the bids are different, i.e., one buyer bids 3 and the other bids 4, then the higher bidder gets the object and pays 3.5.

• **Claim 1.** It is a weakly dominant strategy for both buyers to bid their true valuations.

• **Proof.** Consider Buyer 1; the argument is analogous for Buyer 2.

Suppose his valuation is 3. If Buyer 1 bids 3, then his payoff will be 0 regardless of Buyer 2's bid. If Buyer 1 bids 4, then his payoff will be $3 - 3.5 = -0.5$ if Buyer 2 bids 3 and $0.5 \times (3 - 4) = -0.5$ if Buyer 2 bids 4. Clearly, bidding the true value is better than lying.

Suppose his valuation is 4. If Buyer 1 bids 3, then his payoff will be $0.5 \times (4-3) = 0.5$ if Buyer 2 bids 3, and his payoff will be 0 if Buyer 2 bids 4. If Buyer 1 bids 4, then his payoff will be $4 - 3.5 = 0.5$ if Buyer 2 bids 3, and it will be 0 if Buyer 2 bids 4. Clearly, bidding the true value is no worse than lying.

• **Claim 2.** The seller's expected revenue is 3.5.

• **Proof.** With probability 1/4 both buyers have valuation 3. Since both bid truthfully in equilibrium, the seller's revenue is 3. With probability 1/2, one buyer has valuation 3 and the other has valuation 4. Since both bid truthfully in equilibrium, the seller's revenue is $3.5 = 7/2$. With probability 1/4 both buyers have valuation 4. Since both bid truthfully in equilibrium, the seller's revenue is 4. Thus, the seller's expected revenue is

$$\frac{1}{4} \times 3 + \frac{1}{2} \times \frac{7}{2} + \frac{1}{4} \times 4 = \frac{3}{4} + \frac{7}{4} + \frac{4}{4} = \frac{14}{4} = 3.5$$

• **Claim 3.** The expected payoff of a buyer with valuation 3 is 0. The expected payoff of a buyer with valuation 4 is $0.5 \times (4 - 3.5) = 0.25$. The expected payoff of a buyer prior to knowing his valuation is $0.5 \times 0 + 0.5 \times 0.25 = 0.125$.

• Note that the optimal auction reduces the expected payoff of the high valuation buyer and transfers the difference to the seller. This transfer of rent is similar to what we saw in the case of a discriminating monopolist who extracts the maximum possible rent from the "high type" by pushing down the "low type" to his reservation payoff.

4. Application II: Infinitely repeated games with discounting

4.1 Example: repeated prisoner's dilemma

The following is the well-known Prisoner's Dilemma game.

$$\begin{array}{c} \\ \text{Hawk} \\ \text{Dove} \end{array} \begin{array}{c} \begin{array}{cc} \text{Hawk} & \text{Dove} \end{array} \\ \begin{pmatrix} 0,0 & 3,-1 \\ -1,3 & 2,2 \end{pmatrix} \end{array} \tag{4.1.1}$$

Clearly, both players have a strongly dominant strategy: Hawk. This simple game crystallizes the conflict between private and social welfare.

Suppose the same players play (4.1.1) repeatedly a finite number of times, say n times. The resulting game is called the n-fold repeated game with (4.1.1) as the stage game. Before moving in any stage-game, the players can perfectly observe the history until that stage. We study the case of $n = 2$; analogous results can be easily proved for arbitrary n.

Theorem 4.1.2. *Consider the 2-fold repeated game with (4.1.1) as the stage game. This game has a unique subgame perfect equilibrium, in which both players choose the Scorched Earth strategy: play Hawk in every period regardless of history.*

Proof. Regardless of what happened in the first stage, playing Hawk in Stage 2 dominates playing Dove. Therefore, when pondering the action choice in Stage 1, both players realize that, regardless of what they do in Stage 1, both of them will play Hawk in Stage 2. They understand that, regardless of what they do in Stage 1, their Stage 2 payoff is going to be 0. Therefore, in effect, it is as if they are playing (4.1.1) in Stage 1. In (4.1.1), it is a dominant strategy to choose Hawk. This observation concludes the proof. ■

Remark 4.1.3. *The Scorched Earth strategy is not a dominant strategy in the finitely repeated game; the pair of such strategies is, however, a subgame perfect equilibrium, which is a far less stringent requirement than dominance. For instance, consider an opponent who plays the Grim Trigger strategy: Play Dove in Stage 1; if both players have played Dove in Stage 1, play Dove in Stage 2, otherwise play Hawk in Stage 2. Against the Grim Trigger strategy, the Scorched Earth strategy yields a total payoff of 3, while Grim Trigger yields a total payoff of 4.*

Theorem 4.1.4. *The 2-fold repeated game with (4.1.1) as the stage game has multiple Nash equilibria but all of them lead to the same play: ($(Hawk, Hawk)$ in Stage 1 and $(Hawk, Hawk)$ in Stage 2.*

Proof. Consider the Trust-but-Verify strategy: Play Hawk in Stage 1; if both players plays Dove in Stage 1, play Dove in Stage 2, otherwise play Hawk in Stage 2. The pair of such strategies is a Nash equilibrium, but is not subgame perfect: in the subgame following both players choosing Dove in Stage 1, playing Dove is dominated by playing Hawk.

We now confirm that all Nash equilibria lead to the same play: $(Hawk, Hawk)$ in Stage 1 and $(Hawk, Hawk)$ in Stage 2. Suppose there is a Nash equilibrium that generates a play that involves the players choosing actions other than $(Hawk, Hawk)$ in Stage 2. For the sake of argument, say Player 2 chooses Dove in Stage 2. Clearly, if Player 2 switches his action in Stage 2 to Hawk, then his payoff for the entire game goes up by 1. Thus, Player 2's strategy that lead to the choice of Dove in Stage 2 was not a best response to Player 1's strategy, which contradicts our assumption that the given pair of strategies was a Nash equilibrium.

Thus, every Nash equilibrium of the two-fold repeated Prisoner's dilemma must generate a play that involves playing $(Hawk, Hawk)$ in Stage 2.

Suppose there is a Nash equilibrium that generates a play that involves the players choosing actions other than $(Hawk, Hawk)$ in Stage 1. For the sake of argument, say Player 2 chooses Dove in Stage 1. Clearly, if Player 2 switches his action in Stage 1 to Hawk, then his payoff for the entire game goes up by 1. Thus, Player 2's strategy that lead to the choice of Dove in Stage 1 was not a best response to Player 1's strategy, which contradicts our assumption that the given pair of strategies was a Nash equilibrium.

Thus, every Nash equilibrium of the two-fold repeated Prisoner's dilemma must generate a play that involves playing $(Hawk, Hawk)$ in Stage 1. ■

4.2 General model

Consider the strategic form game $\Gamma = \{N, (C_i, u_i)_{i \in N}\}$. We assume that every C_i is compact subset of $\Re^l$ and every u_i is continuous. Moreover, C_i is always given the discrete σ-algebra. In this section we shall study the game $\Gamma^\infty(\delta)$ generated by repeatedly playing Γ a denumerable number of times. The motivation for studying this game is that it provides a formalism for investigating the establishment and stability of norms of behavior that might be supported by the fact that players have to interact repeatedly and at no stage is the interaction seen as terminal. For instance, it is a commonplace conjecture that the possibility of future retribution in long-run interactions imposes constraints on players' behavior that might not exist in short-run interactions. Infinitely repeated games allow us to model the feasibility and evaluate the potency of using future punishment to induce compliant behavior in the short-run.

There are two broad questions asked in the literature: (a) What payoffs are generated by equilibrium behavior in a repeated game?, and (b) What is the nature of the supporting equilibria? The former is a question regarding the existence of appropriate equilibria, while the latter is a characterization issue.

Given Γ, denote the family of stage games of $\Gamma_\infty(\delta)$ by $\{\Gamma_t \mid t \in \mathcal{N}\}$, where $\Gamma_t = \{N, (C_i^t, u_i)_{i\in N}\}$ with $C_i^t = C_i$ for every $i \in N$ and $t \in \mathcal{N}$. The family of stage games is ordered by the usual order on $\mathcal{N}$. Thus, Γ_t is a copy of Γ with the interpretation that it is played after observing the outcomes of $\{\Gamma_1, \ldots, \Gamma_{t-1}\}$ and prior to the play of Γ_{t+1}. $C^t = \prod_{i\in N} C_i^t$ is the set of action profiles for Γ_t. Set $C^0 = \{h_0\}$. Let $H^t = \prod_{\tau=0}^{t} C^\tau$ and $H = \prod_{t\in \mathcal{Z}_+} C^t$. H is the set of all possible plays of $\Gamma_\infty(\delta)$. Given $h \in H$, h_t is the outcome of Γ_t. Given $h \in H$ (resp. $h^t \in H^t$ and $\tau \leq t$), let $h^\tau = (h_0, \ldots, h_\tau) \in H^\tau$.

For every $T \in \mathcal{N}$, let $\Gamma_\infty^T(\delta)$ be the repeated game with Γ_T as the first stage game. For every $T \in \mathcal{N}$, $\Gamma_\infty^T(\delta)$ is a proper subgame of $\Gamma_\infty(\delta)$.

A strategy for player i in $\Gamma_\infty(\delta)$ is a sequence of functions $s_i = (s_i^t)_{t\in\mathcal{N}}$, where $s_i^t : H^{t-1} \to C_i$. Given $h^{t-1} \in H^{t-1}$, $s_i^t(h^{t-1})$ is the action chosen by player i in Γ_t conditional on the history until that stage. $s^t = (s_i^t)_{i\in N} : H^{t-1} \to C$ is the function that generates the action profile in Γ_t. Let S_i be the set of possible strategies for player i and let $S = \prod_{i\in N} S_i$. When the strategy s_i is understood, we shall abbreviate $s_i^t(h^{t-1})$ to $s_i(h^{t-1})$ and $s^t(h^{t-1})$ to $s(h^{t-1})$. The implementation of strategy $s_i \in S_i$ in the subgame Γ_∞^T amounts to the implementation of the restriction $(s_i^t)_{t=T}^\infty$.

We now describe the play $h(s) \in H$ of $\Gamma_\infty(\delta)$ generated by a strategy profile $s \in S$. Set $h^0(s) = h_0$. Proceeding inductively, given $h^t(s)$, let

$$\sigma^{t+1}(s) = s \circ h^t(s) \quad \text{and} \quad h^{t+1}(s) = (h^t(s), \sigma^{t+1}(s))$$

respectively. Given a strategy profile s, $h^t(s)$ is the history of outcomes until Γ_{t+1}, while $\sigma^t(s)$ is the outcome of Γ_t.

Given $T \in \mathcal{N}$ and a strategy profile $s \in S$, define $U_i^T : S \to \Re$ by

$$U_i^T(s) = (1-\delta) \lim_{\tau \uparrow \infty} \sum_{t=T}^{\tau} \delta^{t-1} u_i \circ \sigma^t(s)$$

$U_i^T(s)$ is player i's payoff in the subgame Γ_∞^T if the strategy profile s is implemented in this subgame. As every C_i is compact, so is C by Tychonoff's theorem. As u_i is continuous, $u_i(C) \subset \Re$ is compact; in particular, $u_i(C)$ is bounded. Therefore, as $\delta \in (0,1)$, the series on the right-hand-side will converge. Now we have the strategic form game $\Gamma_\infty(\delta) = \{N, (S_i, U_i^1)_{i \in N}\}$.

4.3 Nash folk theorem

Given the stage game Γ, define $u : C \to \Re^N$ by $u(c) = (u_i(c))_{i \in N}$. Given u, the set of payoff profiles of Γ is $u(C)$. For player i, define

$$v_i = \min \left\{ \max_{c_i \in C_i} u_i(c_i, c_{-i}) \,\middle|\, c_{-i} \in C_{-i} \right\}$$

v_i is called player i's minmax level; it is the lowest payoff to which player i can be forced by the other players, should they be able to correlate their actions appropriately. Given $v = (v_i)_{i \in N}$, a profile of payoffs $w \in \Re^N$ is enforceable (resp. strictly enforceable) if $w \geq v$ (resp. $w \gg v$) and $w = u(c)$ for some $c \in C$. An outcome $c \in C$ is said to be enforceable (resp. strictly enforceable) if $u(c)$ is enforceable (resp. strictly enforceable). Let $B_i : C_{-i} \rightrightarrows C_i$ be player i's best response mapping and let $b_i : C_{-i} \to C_i$ be a selection from B_i. By definition,

$$v_i = \min_{c_{-i} \in C_{-i}} u_i(b_i(c_{-i}), c_{-i})$$

Let p_{-i} be a solution of this problem, i.e., $u_i(b_i(p_{-i}), p_{-i}) \leq u_i(b_i(c_{-i}), c_{-i})$ for every $c_{-i} \in C_{-i}$. p_{-i} represents the severest punishment that the other players can inflict upon player i. Let $p^i = (b_i(p_{-i}), p_{-i})$ be the outcome when player i's rivals minmax player i and player i responds optimally to their choices.

The folk theorems provide links between the set of equilibrium outcomes of $\Gamma_\infty(\delta)$ and $u(C)$. As we vary the notion of equilibrium, we get various folk theorems associated with these equilibrium concepts. The following result shows that every Nash equilibrium payoff profile of the repeated game is an enforceable profile for the stage game.

Theorem 4.3.1. *If $\delta \in (0,1)$ and s is a Nash equilibrium of $\Gamma_\infty(\delta)$, then the payoff profile $U(s)$ is an enforceable profile for the stage game Γ.*

Proof. Suppose $\delta \in (0,1)$, s is a Nash equilibrium of $\Gamma_\infty(\delta)$ and $U(s)$ is not an enforceable profile for Γ. As $U(s)$ is not enforceable, there exists $i \in N$ such that $U_i(s) < v_i$. We construct a strategy $\bar{s}_i$ for player i in $\Gamma_\infty(\delta)$ such that $U_i(s_{-i}, \bar{s}_i) \geq v_i$. Consequently, $U_i(s_{-i}, \bar{s}_i) > U_i(s)$, a contradiction.

Let $\bar{s}_i^j(h) = b_i \circ s_{-i}^j(h)$ for every $h \in H^j$ and every $j \in \mathcal{N}$. By definition,

$$u_i(s_{-i}^j(h), \bar{s}_i^j(h)) = u_i(s_{-i}^j(h), b_i \circ s_{-i}^j(h)) \geq u_i(p_{-i}, b_i(p_{-i})) = v_i$$

for every $h \in H^j$ and every $j \in \mathcal{N}$. Let $\bar{h}^j \in H^j$ be the history until Γ_j generated by the strategy profile $(s_{-i}, \bar{s}_i)$, i.e., $\bar{h}^j = h^j(s_{-i}, \bar{s}_i)$. Thus,

$$u_i \circ \sigma^j(s_{-i}, \bar{s}_i) = u_i(s_{-i}^j(\bar{h}^j), \bar{s}_i^j(\bar{h}^j)) \geq v_i$$

for every $j \in \mathcal{N}$. Consequently,

$$(1-\delta)\sum_{j=1}^{J} \delta^{j-1} u_i \circ \sigma^j(s_{-i}, \bar{s}_i) \geq (1-\delta)\sum_{j=1}^{J} \delta^{j-1} v_i = v_i(1-\delta)\sum_{j=1}^{J} \delta^{j-1}$$

for every $J \in \mathcal{N}$. Letting $J \uparrow \infty$, we have $U_i(s_{-i}, \bar{s}_i) \geq v_i$. ■

We now prove an approximate converse of this result.

Theorem 4.3.2. *If w is a strictly enforceable payoff profile of Γ, then there exists $\underline{\delta} \in [0,1)$ such that, for every $\delta \in (\underline{\delta}, 1)$, there exists a Nash equilibrium s of $\Gamma_\infty(\delta)$ such that $U(s) = w$.*

Proof. Suppose w is a strictly enforceable payoff of Γ. Thus, $w \gg v$ and there exists $c \in C$ such that $u(c) = w$. We define player i's strategy for $\Gamma_\infty(\delta)$ as follows. Let $s_i^1(h^0) = c_i$. Suppose the functions $\{s^1, \ldots, s^t\}$ are given. Consider Γ_{t+1} and $h^t \in H^t$. Given $\tau \in \{1, \ldots, t\}$, let[10]

$$n(h^t, \tau) = |\{j \in N \mid s_j^\tau \circ h^\tau \neq c_j\}| \quad \text{and} \quad T = \inf\{\tau \in \{1, \ldots, t\} \mid n(h, \tau) = 1\}$$

[10] By convention, the infimum of an empty set is ∞.

$n(h,\tau)$ is the number of players $j \in N$ who chose actions other than c_j in Γ_τ. T is the first time when exactly one player j deviates from c_j. Then,

$$s_i^{t+1}(h^t) = \begin{cases} c_i, & \text{if } t < T \\ p_i^j, & \text{if } t \geq T \text{ and } c_j \neq s_j^T(h^T) \end{cases}$$

Player i's strategy requires i to play c_i in every stage game until in some stage game exactly one player, say player j, deviates from playing c_j. In every subsequent stage game, player i is required to 'punish' player j by playing p_i^j.[11] As punishment is triggered by the deviation of a single player j, the punishment p^j is well-defined. We show that this profile of strategies is a Nash equilibrium.

Clearly, $h(s) = (c, c, \ldots)$. Suppose player i uses strategy s_i' instead of s_i. Let

$$T^* = \inf\{t \in \mathcal{N} \mid h_t(s_{-i}, s_i') \neq h_t(s)\}$$

If $T^* = \infty$, then $U_i(s_{-i}, s_i') = U_i(s)$. Suppose $T^* \in \mathcal{N}$. By definition, $h^{T^*-1}(s_{-i}, s_i') = (c, \ldots, c)$. Then, $s_{-i} \circ h^{T^*-1}(s_{-i}, s_i') = c_{-i}$ and $s_i \circ h^{T^*-1}(s_{-i}, s_i') \neq c_i$. Consequently, $s_{-i} \circ h^{T^*-1+k}(s_{-i}, s_i') = p_{-i}^i$ for every $k \in \mathcal{N}$. Therefore, setting $r_i = u_i(b_i(c_{-i}), c_{-i})$, we have $r_i \geq w_i > v_i$ and

$$\begin{aligned} \frac{U_i(s_{-i}, s_i')}{1-\delta} &\leq \sum_{t=1}^{T^*-1} \delta^{t-1} u_i(c) + \delta^{T^*-1} u_i(s_i' \circ h^{T^*-1}(s_{-i}, s_i'), c_{-i}) + \frac{\delta^{T^*}}{1-\delta} u_i(p^i) \\ &\leq \sum_{t=1}^{T^*-1} \delta^{t-1} w_i + \delta^{T^*-1} u_i(b_i(c_{-i}), c_{-i}) + \frac{\delta^{T^*}}{1-\delta} v_i \\ &= \sum_{t=1}^{T^*-1} \delta^{t-1} w_i + \delta^{T^*-1} r_i + \frac{\delta^{T^*}}{1-\delta} v_i \end{aligned}$$

The first inequality follows from the fact that $u_i(p^i) \geq u_i(s_i' \circ h^t(s_{-i}, s_i'), p_{-i})$ for every $t \geq T^*$. Also,

$$\frac{U_i(s)}{1-\delta} = \sum_{t=1}^{T^*-1} \delta^{t-1} w_i + \delta^{T^*-1} w_i + \frac{\delta^{T^*}}{1-\delta} w_i$$

[11] This strategy is called the "grim trigger". Note that, conditional on punishment being triggered, the subsequent strategy for a punisher is an open-loop strategy, i.e., it is history-independent *modulo* the triggering event.

Then,

$$\frac{U_i(s) - U_i(s_{-i}, s_i')}{1-\delta} \geq \delta^{T^*-1}(w_i - r_i) + \frac{\delta^{T^*}}{1-\delta}(w_i - v_i)$$
$$= \delta^{T^*-1}\left[(w_i - r_i) + \frac{\delta}{1-\delta}(w_i - v_i)\right]$$

Setting $\underline{\delta} = (r_i - w_i)/(r_i - v_i) \in [0, 1)$, we have $U_i(s) \geq U_i(s_{-i}, s_i')$ for every $\delta \in (\underline{\delta}, 1)$. ■

4.4 Subgame perfect folk theorem

The following result is an example of a subgame perfect folk theorem. While it is not as general or subtle in its construction as many later subgame perfect folk theorems, it serves to illustrate the steps required in proving such a result.

Theorem 4.4.1. *(Friedman) If e is a Nash equilibrium of Γ and w is an enforceable payoff profile such that $w \gg u(e)$, then there exists $\underline{\delta} \in [0, 1)$ such that, for every $\delta \in (\underline{\delta}, 1)$, there exists a subgame perfect equilibrium s of $\Gamma_\infty(\delta)$ such that $U(s) = w$.*

Proof. Suppose e is a Nash equilibrium of Γ and $w \gg u(e)$. By definition, there exists $c \in C$ such that $u(c) = w$. We define player i's strategy for $\Gamma_\infty(\delta)$ as follows. Let $s_i(h^0) = c_i$. Suppose $s_i(h^{t-1})$ is given for every $h^{t-1} \in H^{t-1}$ and every $t \in \mathcal{N}$. Consider Γ_{t+1} and $h^t \in H^t$. Then,

$$s_i(h^t) = \begin{cases} c_i, & \text{if } h^t = (c, \ldots, c) \\ e_i, & \text{otherwise} \end{cases}$$

Player i's strategy requires i to play c_i in every stage game until in some stage game a deviation from the profile c is observed. In every subsequent stage game, player i is required to play his part e_i of the stage-game Nash equilibrium e, which plays the role of a collective punishment.[12] We show that this profile of strategies is a subgame perfect equilibrium.

Fix $T \in \mathcal{N}$ and $h^{T-1} \in H^{T-1}$. This fixes the proper subgame $\Gamma_\infty^T(\delta)$ of $\Gamma_\infty(\delta)$. Suppose player i implements strategy s_i' instead of s_i in $\Gamma_\infty^T(\delta)$. Let

$$\eta = \inf\{t \in \{1, \ldots, T-1\} \mid h_t \neq c\} \quad \text{and} \quad \tau = \inf\{t \in \mathcal{N} \mid t \geq T \;\wedge\; h_t(s_{-i}, s_i') \neq h_t(s)\}$$

[12] This strategy is called "Nash reversion". Note that, conditional on collective punishment being triggered, the subsequent strategy for every player is an open-loop strategy, i.e., it is history-independent *modulo* the triggering event.

Suppose $\eta = \infty$, i.e., $h^{T-1} = (c, \ldots, c)$, which means no deviation from c is observed in the history prior to the subgame $\Gamma^T_\infty(\delta)$. If $\tau = \infty$, then $h_t(s_{-i}, s'_i) = h_t(s)$ for every $t \geq T$. Consequently, $U^T_i(s_{-i}, s'_i) = U^T_i(s)$. Alternatively, suppose $\tau \in \mathcal{N}$. By definition, $h^{\tau-1}(s_{-i}, s'_i) = h^{\tau-1}(s) = (c, \ldots, c)$. Then, $s_{-i} \circ h^{\tau-1}(s_{-i}, s'_i) = c_{-i}$ and $s'_i \circ h^{\tau-1}(s_{-i}, s'_i) \neq c_i$. Consequently, $s_{-i} \circ h^{\tau-1+k}(s_{-i}, s'_i) = e_{-i}$ for every $k \in \mathcal{N}$. Therefore, setting $r_i = u_i(c_{-i}, b_i(c_{-i}))$, we have $r_i \geq w_i$ and

$$\begin{aligned}
\frac{U^T_i(s_{-i}, s'_i)}{1-\delta} &\leq \sum_{t=T}^{\tau} \delta^{t-1} u_i \circ (s_{-i}, s'_i) \circ h^{t-1}(s_{-i}, s'_i) + \frac{\delta^\tau}{1-\delta} u_i(e) \\
&\leq \sum_{t=T}^{\tau-1} \delta^{t-1} u_i(c) + \delta^{\tau-1} u_i(c_{-i}, b_i(c_{-i})) + \frac{\delta^\tau}{1-\delta} u_i(e) \\
&= \sum_{t=T}^{\tau-1} \delta^{t-1} w_i + \delta^{\tau-1} r_i + \frac{\delta^\tau}{1-\delta} u_i(e)
\end{aligned}$$

The first inequality follows from the fact that $u_i(e) \geq u_i(s'_i \circ h^t(s_{-i}, s'_i), e_{-i})$ for every $t \geq \tau$. Also,

$$\frac{U^T_i(s)}{1-\delta} = \sum_{t=T}^{\tau-1} \delta^{t-1} w_i + \delta^{\tau-1} w_i + \frac{\delta^\tau}{1-\delta} w_i$$

Then,

$$\begin{aligned}
\frac{U^T_i(s) - U^T_i(s_{-i}, s'_i)}{1-\delta} &\geq \delta^{\tau-1}(w_i - r_i) + \frac{\delta^\tau}{1-\delta}[w_i - u_i(e)] \\
&= \delta^{\tau-1}\left[(w_i - r_i) + \frac{\delta}{1-\delta}[w_i - u_i(e)]\right]
\end{aligned}$$

Setting $\underline{\delta} = (r_i - w_i)/[r_i - u_i(e)] \in [0, 1)$, we have $U^T_i(s) \geq U^T_i(s_{-i}, s'_i)$ for every $\delta > \underline{\delta}$.

Suppose $\eta \in \{1, \ldots, T-1\}$. Then, $h^{T-1} \neq (c, \ldots, c)$, and consequently, $s_{-i} \circ h^{t-1}(s_{-i}, s'_i) = e_{-i} = s_{-i} \circ h^{t-1}(s)$ for every $t \in \mathcal{N}$ such that $t \geq T$. Therefore,

$$\begin{aligned}
\frac{U^T_i(s_{-i}, s'_i)}{1-\delta} &= \sum_{t=T}^{\infty} \delta^{t-1} u_i \circ (s_{-i}, s'_i) \circ h^{t-1}(s_{-i}, s'_i) \\
&\leq \sum_{t=T}^{\infty} \delta^{t-1} u_i(e_{-i}, b_i(e_{-i})) = \sum_{t=T}^{\infty} \delta^{t-1} u_i(e) = \frac{U^T_i(s)}{1-\delta}
\end{aligned}$$

as required. ■

References

Berge, C. (1963), *Topological Spaces*, Macmillan, New York.

Browder, F. E. (1968), The fixed point theory of multi-valued mappings in topological vector spaces, *Mathematische Annalen* **177**, 283-301.

Fan, K. (1952), Fixed-point and minimax theorems in locally convex topological linear spaces, *Proceedings of the National Academy of Sciences* **38**, 121-126.

Harsanyi, J. C. (1967/68), Games with incomplete information played by Bayesian players, Parts I, II and III, *Management Science* **14**, 159-182, 320-334 and 486-502.

Herstein and Milnor

Kakutani, S. (1941), A generalization of Brouwer's fixed point theorem, *Duke Mathematical Journal* **8**, 457-59.

Kuhn, H. W. (1953), Extensive games and the problem of information, pp. 193-216, in *Contributions to the Theory of Games, Volume II*, Annals of Mathematical Studies, 28, H. W. Kuhn and A. W. Tucker, eds., Princeton University Press, Princeton.

Mertens, J.-F. and S. Zamir (1985), Formulation of Bayesian analysis for games with incomplete information, *International Journal of Game Theory* **14**, 1-29.

Nash, J. F. (1950), Equilibrium points in N-person games, *Proceedings of the National Academy of Sciences of the United States of America* **36**, 48-49.

Selten, R. (1975), Reexamination of the perfectness concept for equilibrium points in extensive games, *International Journal of Game Theory* **4**, 25-55.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*, John Wiley and Sons, New York.