Conditional strategy equilibrium*

Lorenzo Bastianello[†] Mehmet S. Ismail[‡]

March 11, 2021

1 Definition of pure conditional strategies

Let $(A_i, u_i)_{i\in N}$ be an n-person noncooperative game, where $N = \{1, ..., n\}$ is the finite set of players, $A_i = \{a_i^1, a_i^2, ..., a_i^{m_i}\}$ finite nonempty action set, and $u_i : A \to \mathbb{R}_+$ the (von Neumann-Morgenstern) utility function of player $i \in N$. A strategy profile is denoted by $a \in A = \times_{i\in N} A_i$.[1]

A pure conditional strategy of player i is a function $s_i : A_{-i} \to A_i$. A pure conditional strategy profile is $s \in S = \times_{i\in N} S_i$ where S_i is the set of conditional strategies of player i. A pure conditional strategy profile $s \in S$ is an *agreement* if there exists a unique action profile $a \in A$ such that $s(a) = a$—i.e., for all i $s_i(a_{-i}) = a_i$—which is the unique fixed point of the conditional strategy profile s, which always exists because for each i A_i is nonempty. The intuition is that everybody agrees on the unique strategy profile a. If $s \in S$ is not an agreement, then it is called a *disagreement*.

Let $B = \{s \in S | s \text{ is an agreement}\}$ be the set of agreements. The utility function $U_i : S \to \mathbb{R}_+$ is defined as follows.

$$U_i(s) = \begin{cases} u_i(a) & \text{if } s \in B \\ 0 & \text{if } s \in S \setminus B, \end{cases}$$

where $a = s(a)$. This is a well-defined function since every $s \in B$ has a unique fixed point, and the disagreement payoff is zero. Note that $U_i|_B = u_i$.

*This is a preliminary draft. The prior literature on conditional actions, starting with (Nigel, 1971), will be reviewed in the next version.

[†]LEMMA, Université Paris 2 Panthéon-Assas, Paris, France. E-mail: lorenzo.bastianello@u-paris2.fr

[‡]Department of Political Economy, King's College London, UK. E-mail: mehmet.s.ismail@gmail.com.

[1]Without loss of generality, we consider nonnegative utility function. The number of pure conditional actions of player i is given by $m_i^{\prod_{j\neq i} m_j} = |S_i|$ because $\prod_{j\neq i} m_j$ gives the number of pure action profiles of everyone but i.

Let $(S_i, U_i)_{i\in N}$ denote a game in *conditional extension* of the game $(A_i, u_i)_{i\in N}$. We next define conditional strategy equilibrium.

Definition 1. *A conditional strategy profile $s^* \in S$ is called a conditional strategy equilibrium if for all i $s_i^* \in \arg\max U_i(s^*)$ or equivalently*

$$U_i(s^*) \geq U_i(s_i', s_{-i}^*).$$

In other words, a conditional strategy equilibrium (CSE) is a self-enforcing agreement in the space of conditional strategy profiles, just like a Nash equilibrium (Nash, 1950) is a self-enforcing agreement in the space of mixed strategy profiles.

2 Conditional mixed extension

Let $G = (\Delta A_i, u_i)_{i\in N}$ be an n-person noncooperative game in mixed extension, where $N = \{1, ..., n\}$ is the finite set of players, ΔA_i the set of all probability distributions over the finite action set A_i, and $u_i : \Delta A \to \mathbb{R}$ the von Neumann-Morgenstern expected utility function of player $i \in N$ (von Neumann and Morgenstern, 1944). A mixed strategy profile is denoted by $p \in \Delta A = \times_{i\in N}\Delta A_i$, where ΔA_i is the simplex in $\mathbb{R}^{m_i-1}$.

Let $\mathcal{A}_{-i}$ be the Borel σ-algebra over ΔA_{-i}. The set of conditional mixed strategies of player i is given by $\Sigma_i = \{\ \sigma_i : \Delta A_{-i} \to \Delta A_i \mid \sigma_i \text{ is a simple and } \mathcal{A}_{-i}\text{-measurable}\}$. A conditional mixed strategy profile is $\sigma \in \Sigma = \times_{i\in N}\Sigma_i$ where Σ_i is the set of conditional strategies of player i. A conditional strategy profile $\sigma \in \Sigma$ is an *agreement* if there exists a unique $p \in \Delta A$ such that $\sigma(p) = p$—i.e., for all i $\sigma_i(p_{-i}) = p_i$—which is the unique fixed point of the conditional strategy profile σ, which always exists because for each i ΔA_i is nonempty. The intuition is that everybody agrees on the unique strategy profile p. If $\sigma \in \Sigma$ is not an agreement, then it is called a *disagreement*.

Let $B = \{\sigma \in \Sigma | \sigma \text{ is an agreement}\}$ be the set of agreements. The utility function $U_i : \Sigma \to \mathbb{R}_+$ is defined as follows.

$$U_i(\sigma) = \begin{cases} u_i(p) & \text{if } \sigma \in B \\ 0 & \text{if } \sigma \in \Sigma \setminus B, \end{cases}$$

where $p = \sigma(p)$. This is a well-defined function since every $\sigma \in B$ has a unique fixed point. Note that we normalize the disagreement payoff to zero and that $U_i|_B = u_i$.

Let $\Gamma = (\Sigma_i, U_i)_{i \in N}$ denote a game in *conditional mixed extension* of the game $(A_i, u_i)_{i \in N}$. We next define conditional strategy equilibrium.

Definition 2. *A conditional strategy profile $\sigma^* \in \Sigma$ is called a conditional strategy equilibrium if for all i $\sigma_i^* \in \arg\max U_i(\sigma^*)$ or equivalently*

$$U_i(\sigma^*) \geq U_i(\sigma_i', \sigma_{-i}^*).$$

In other words, a conditional mixed strategy equilibrium (CSE) is a self-enforcing agreement in the space of conditional strategies.

2.1 Conditional mixed extension and the mixed extension of the pure conditional strategies

Consider $\hat{S}_i = \{\hat{s}_i : \Delta A_{-i} \to A_i \mid \hat{s}_i \text{ is } \mathcal{A}_{-i}\text{-measurable}\}$, which is the set of pure conditional strategies of player i against mixed strategies of the others. We denote $\Delta \hat{S}_i$ the set of all probability measures with finite support over $\hat{S}_i$. Every probability measure $\mu \in \Delta \hat{S}_i$ induces a function $\sigma_i \in \Sigma_i$. Define function

$$\begin{array}{rccc} \phi : & \Delta \hat{S}_i & \to & \Sigma_i \\ & \mu & \mapsto & \sigma_i^\mu \end{array} \tag{1}$$

where for all q_{-i}

$$\sigma_i^\mu(q_{-i}) = \int_{\hat{S}_i} \hat{s}_i(q_{-i}) d\mu(\hat{s}_i),$$

which is a probability distribution over A_i given q_{-i}.

Next, we show that every conditional mixed strategy σ_i can be induced by a probability measure μ over conditional (pure) strategies.

Theorem 1. *For every $\sigma_i \in \Sigma_i$ there exists a $\mu \in \Delta \hat{S}_i$ such that $\phi(\mu) = \sigma_i$ where ϕ is defined in (1).*

Proof. Fix $\sigma_i \in \Sigma_i$ and consider the induced partition over ΔA_{-i} $\mathcal{P}_{\sigma_i} = \{X_1, ..., X_{L_{\sigma_i}}\}$. For ease of notation, we denote $\mathcal{P} = \mathcal{P}_{\sigma_i}$ and $L = L_{\sigma_i}$. Note that σ_i can be written as

$$\begin{array}{rccc} \sigma_i : & \mathcal{P} & \to & \Delta A_i \\ & l & \mapsto & \sigma_i(l) = \mu_l \end{array}$$

where μ_l is a probability measure over A_i. Put differently, for $a_i \in A_i$, $\mu_l(a_i) = \sigma_i(q_{-i})(a_i)$ for all $q_{-i} \in X_l$. The finite collection of probability spaces $(A_i, 2^{A_i}, \mu_l)_{l=1}^L$ induces the product space $(\times_{l=1}^{L} A_i, 2^{\times_{l=1}^{L} A_i}, \mu)$ where μ is the usual product measure over $2^{\times_{l=1}^{L} A_i}$ defined by $\mu((a_i^1, a_i^2, ..., a_i^L)) = \prod_{l=1}^{L} \mu_l(a_i^l)$ for all $(a_i^1, a_i^2, ..., a_i^L)$ in $\times_{l=1}^{L} A_i = \{(a_i^1, a_i^2, ..., a_i^L) | a_i^l \in A_i\}$.

Note that every element $(a_i^1, a_i^2, ..., a_i^L) \in \times_{l=1}^{L} A_i$ can be identified with a $\hat{s}_i \in \hat{S}_i|_{\mathcal{P}}$ where $\hat{S}_i|_{\mathcal{P}} = \{\hat{s}_i \in \hat{S}_i | \hat{s}_i \text{ is } \mathcal{P}\text{-measurable}\}$, i.e., $\hat{s}_i(q_{-i}) = a_i^l$ for all l and all $q_{-i} \in X_l$. Thus, μ is a measure supported by $\hat{S}_i|_{\mathcal{P}}$ such that for all $\hat{s}_i \in \hat{S}_i|_{\mathcal{P}}$, $\mu(\hat{s}_i) = \prod_{l=1}^{L} \mu_l(\hat{s}_i(l))$ where $\hat{s}_i(l) = \hat{s}_i(q_{-i})$ for all $q_{-i} \in X_l$.

Next we show that $\phi(\mu) = \sigma_i$. Recall that by (1) $\phi(\mu) = \sigma_i^{\mu}$, which is defined by for all q_{-i} and all a_i

$$\sigma_i^{\mu}(q_{-i})(a_i) = \int_{\hat{S}_i} \hat{s}_i(q_{-i})(a_i) d\mu(\hat{s}_i) = \int_{\hat{S}_i|_{\mathcal{P}}} \hat{s}_i(q_{-i})(a_i) d\mu(\hat{s}_i).$$

Note that

$$\hat{s}_i(q_{-i})(a_i) = \begin{cases} 0 \text{ if } \hat{s}_i(q_{-i}) \neq a_i \\ 1 \text{ if } \hat{s}_i(q_{-i}) = a_i, \end{cases}$$

and, therefore,

$$\sigma_i^{\mu}(q_{-i})(a_i) = \int \mathbb{1}_{\{\hat{s}_i \in \hat{S}_i|_{\mathcal{P}} : \hat{s}_i(q_{-i}) = a_i\}} d\mu(\hat{s}_i) = \mu(\{\hat{s}_i \in \hat{S}_i|_{\mathcal{P}} : \hat{s}_i(q_{-i}) = a_i\}).$$

Fix $a_i \in A_i$ and $q_{-i} \in X_{l'}$ for some $l' \in \{1, ..., L\}$. Then, $\{\hat{s}_i \in \hat{S}_i|_{\mathcal{P}} : \hat{s}_i(q_{-i}) = a_i\} = \{\times_{l=1}^{l'-1} A_i \times \{a_i\} \times_{l=l'+1}^{L} A_i\}$ and since μ is the product measure,

$$\mu(\times_{l=1}^{l'-1} A_i \times \{a_i\} \times_{l=l'+1}^{L} A_i) = \mu_{l'}(a_i) = \sigma_i(q_{-i})(a_i).$$

This implies that for all $q_{-i} \in \Delta A_{-i}$ and all $a_i \in A_i$, $\sigma_i^{\mu}(q_{-i})(a_i) = \sigma_i(q_{-i})(a_i)$, i.e., $\phi(\mu) = \sigma_i$. □

Theorem 2. *In every n-person game, there exists a pure conditional equilibrium.*

Proof. Let $BR_i(\sigma_{-i}) = \{\sigma_i \in \Sigma_i | \sigma_i \in \arg\max_{\sigma_i' \in \Sigma_i} U_i(\sigma_i', \sigma_{-i})\}$. Let $\bar{S}_i$ denote the set of constant strategies of player i—i.e., $\bar{s}_i \in \bar{S}_i$ if and only if for every a_{-i} $\bar{s}_i(a_{-i}) = a_i$.

Player 1 chooses $\hat{s}_1$ such that $\hat{s}_1(a_{-1}) \in BR_1(a_{-1})$ for every $a_{-1} \in A_{-1}$. Let $\hat{a}_2(a_3, ..., a_n) \in \arg\max_{a_2' \in A_2} u_2(\hat{s}_1(a_2', a_3, ..., a_n), a_2', a_3, ..., a_n)$. Then, define the conditional strategy of player 2 as $\hat{s}_2(a_{-2}) = \hat{a}_2(a_3, ..., a_n)$ for every a_{-2}. In words, fixing the action profile $(a_3, ..., a_n)$ player 2 chooses an action that maximizes her utility given that player 1 best responds to her action.

Analogously, consider $\hat{a}_3(a_4, ..., a_n)$ in

$$\arg\max_{a'_3 \in A_3} u_3(\underbrace{\hat{s}_1(\underbrace{\hat{a}_2(a'_3, ..., a_n)}_{\hat{a}_2}, a'_3, a_4, ..., a_n)}_{\hat{a}_1}, \underbrace{\hat{a}_2(a'_3, ..., a_n)}_{\hat{a}_2}, a'_3, a_4, ..., a_n).$$

Then, define the conditional strategy of player 3 as $\hat{s}_3(a_{-3}) = \hat{a}_3(a_4, ..., a_n)$ for every a_{-3}. Conditional strategy of player k is defined as $\hat{s}_k(a_{-k}) = \hat{a}_k(a_{(k+1)}, ..., a_n)$ for every a_{-k} where $\hat{a}_k$ is defined analogously. For player n, $\hat{a}_n \in \arg\max_{a'_n \in A_n} u_n(\hat{a}_1, ..., \hat{a}_{n-1}, a'_n)$, and $\hat{s}(a_{-n}) = \hat{a}_n$. Define with a slight abuse of notation $\hat{a} = (\hat{a}_1, \hat{a}_2, ..., \hat{a}_n)$ where $\hat{a}_i = \hat{a}_i(\hat{a}_{i+1}, ..., \hat{a}_n)$

Next, we show that $\hat{s} \in B$, i.e., it is an agreement. We will prove that $s(\hat{a}) = \hat{a}$ and that for all a such that $\hat{s}(a) = a$, then $a = \hat{a}$. This is because $\hat{s}_n(a_{-n}) = \hat{a}_n$ by definition of $\hat{s}_n$, so $a_n = \hat{a}_n$. Then, given $\hat{a}_n$, $\hat{s}_{n-1}(a_{-(n-1)}) = \hat{a}_{n-1}(a_n) = \hat{a}_{n-1}(\hat{a}_n) = \hat{a}_{n-1}$ by definition of $\hat{s}_n$, so $a_{n-1} = \hat{a}_{n-1}$. We repeat this process till we obtain $a_2 = \hat{a}_2$. Then, notice that $\hat{s}_1(a_{-1}) = BR_1(a_{-1}) = BR_1(\hat{a}_{-1}) = \hat{a}_1$.

Now we show that $\hat{s}$ is a conditional equilibrium. First, we prove that player n has no unilateral profitable deviation from $\hat{s}$. Suppose by way of contradiction that $\tilde{s}_n$ is such a deviation. Consider $\tilde{s} = (\tilde{s}_n, \hat{s}_{-n})$, which must be an agreement in B because otherwise $U_n(\tilde{s}) = 0$. Then, there must be a unique $\tilde{a} \in A$ such that $\tilde{s}(\tilde{a}) = \tilde{a}$. Let $\tilde{a}_n = \tilde{s}_n(\tilde{a}_{-n})$. Then, $u_n(\tilde{a}_n, \hat{a}_{-n}) \leq u_n(\hat{a}) = \max_{a'_n \in A_n} u_n(a'_n, \hat{a}_{-n})$. Thus, $\tilde{s}_n$ is not a unilateral profitable deviation. For every player $k > 1$, the reasoning is analogous. Finally, we show that player 1 has no unilateral profitable deviation. Fixing $\hat{s}_{-1}$ induces the action profile $\hat{a}_{-1}$. Since $\hat{s}_1$ is by definition a best reply to $\hat{a}_{-1}$, player 1 cannot strictly benefit from a unilateral deviation. $\square$

References

Nash, J. F. (1950). Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences 36*(1), 48–49.

Nigel, H. (1971). Paradoxes of Rationality: Theory of Metagames and Political Behavior.

von Neumann, J. and O. Morgenstern (1944). *Theory of Games and Economic Behavior* (1953, 3rd ed.). Princeton: Princeton University Press.