

Correlated Quantal Responses and Equilibrium Selection*

Takashi Ui[†]
Faculty of Economics
Yokohama National University
oui@ynu.ac.jp

February 2001

Revised, October 2004

Abstract

This paper considers incomplete information games with payoffs subject to correlated random disturbances. It explains the connection between the uniqueness of quantal response equilibria, where *large* noise is required, and the uniqueness of equilibria in global games, where *small* noise is required.

JEL classifications: C72, D82.

Keywords: global game; quantal response equilibrium.

1 Introduction

This paper considers a class of incomplete information games with payoffs subject to correlated random disturbances in order to give a unified view on the following two distinct results on the uniqueness of equilibria.

One is the uniqueness of quantal response equilibria á la McKelvey and Palfrey (1995). They considered a model in which players make choices based on quantal response models and defined a quantal response equilibrium. The model is equivalent to an incomplete information game where the actual payoff structure is determined by payoffs of some fixed game plus independent random terms and each player's private signal is his own payoffs. A quantal response equilibrium is a probability distribution of action profiles in a Bayesian Nash equilibrium. McKelvey and Palfrey (1995) showed that if the variance of the independent random terms is large enough then a quantal response equilibrium is unique.

The other is the uniqueness of equilibria in global games á la Carlsson and van Damme (1993). A global game is an incomplete information game where the actual payoff structure is determined by a random draw from a given class of games and each player makes an

*I am grateful for valuable discussions with Atsushi Kajii and Stephen Morris. I acknowledge Financial Support from the Zengin Foundation for Studies on Economics.

[†]*Corresponding author's address:* Faculty of Economics, Yokohama National University, 79-3 Tokiwadai, Hodogaya-ku, Yokohama 240-8501, JAPAN. Phone: +81-45-339-3531. Fax: +81-45-339-3574.

independent noisy observation of the selected game. Different from the game of McKelvey and Palfrey (1995), each player's private signal is not his own payoffs, and these signals are correlated. Carlsson and van Damme (1993) showed that if the variance of the independent noises in the observation is small enough then a Bayesian Nash equilibrium is unique.

The former uniqueness comes from the *large* variance and the latter uniqueness comes from the *small* variance. The natural question is how these different results are related.

Motivated by this question, we introduce correlations of quantal responses in the game of McKelvey and Palfrey (1995). That is, we consider an incomplete information game where the actual payoff structure is determined by payoffs of some fixed 2×2 game plus *correlated* random terms and each player's private signal is his own payoffs. We assume that private signals are jointly normally distributed, and parametrize the class of incomplete information games in terms of the variances and the correlation coefficients. Using the technique of normally distributed global games (Carlsson and van Damme, 1993; Morris and Shin, 2002), we provide a sufficient condition for the uniqueness of Bayesian Nash equilibria. The condition states that if the variance is larger than a threshold given by a decreasing function of the correlation coefficient, then a Bayesian Nash equilibrium is unique. Thus, for the equilibrium to be unique, the variance must be large if the correlation is low, which corresponds to the uniqueness of quantal response equilibria, and the correlation must be high if the variance is small, which corresponds to the uniqueness of equilibria in global games.

One very natural reason for the correlations of quantal responses is that the random terms consist of a common component and an idiosyncratic component. The special case of our games with this interpretation was discussed in Carlsson and van Damme (1993, Appendix B). They demonstrated how the global game analysis and Harsanyi's purification result (Harsanyi, 1973) could be understood as special cases of the same class of games: a very small idiosyncratic component with a fixed common component induces the global game analysis, and a very small idiosyncratic component with no common component induces the purification result.

In contrast, this paper demonstrates how the two distinct *uniqueness* results can be understood as special cases of the same class of games: the distinction is attributed to the large variance with the low correlation and the high correlation with the small variance. Combining our result and Carlsson and van Damme (1993, Appendix B), we can understand global games, quantal response equilibria, and purification within the same framework. This issue is further elaborated by Morris and Shin (2005) in more general setup.

This paper is organized as follows. Section 2 introduces the class of incomplete information games. Section 3 provides the result. Section 4 compares our result and the global games analysis.

2 Model

Consider the following 2×2 symmetric game with the set of players $N = \{1, 2\}$ and the sets of actions $A_1 = A_2 = \{\alpha, \beta\}$.

	α	β
α	x, x	v, u
β	u, v	y, y

We assume that $x - u > 0$ and $y - v > 0$ and thus (α, α) and (β, β) are pure-action Nash equilibria. Players' payoff functions are denoted by g_1 and g_2 . The actual payoffs are subject to random disturbances, $\varepsilon_1 = (\varepsilon_1(\alpha), \varepsilon_1(\beta))$ and $\varepsilon_2 = (\varepsilon_2(\alpha), \varepsilon_2(\beta))$, of which distributions are known to all the players. When ε_i is realized, player i 's payoffs are given by $g_i(a_i, a_j) + \varepsilon_i(a_i)$ for $a_i \in A_i$ and $a_j \in A_j$ with $i \neq j$. Mckelvey and Palfrey (1995) regarded $g_i(a_i, a_j) + \varepsilon_i(a_i)$ as a quantal choice model of player i , assuming that ε_1 and ε_2 are independent, whereas we allow correlations. The realized payoffs are given by the following table.

	α	β
α	$x + \varepsilon_1(\alpha), x + \varepsilon_2(\alpha)$	$v + \varepsilon_1(\alpha), u + \varepsilon_2(\beta)$
β	$u + \varepsilon_1(\beta), v + \varepsilon_2(\alpha)$	$y + \varepsilon_1(\beta), y + \varepsilon_2(\beta)$

Player i 's private signal is his own payoffs, or equivalently, ε_i . We regard this setup as an incomplete information game.

A strategy of player i is a mapping $\sigma_i : \mathbb{R}^2 \rightarrow \Delta(A_i)$. Given σ_i and the private signal ε_i , player i chooses a mixed action $\sigma_i(\varepsilon_i) \in \Delta(A_i)$. We write $\sigma_i(a_i|\varepsilon_i)$ for the probability of choosing $a_i \in A_i$. We write Σ_i for the set of strategies of player i , and write $\Sigma = \Sigma_1 \times \Sigma_2$.

Let $E[\cdot|\varepsilon_i]$ be a conditional expectation operator given ε_i . Then $E[\sigma_j(a_j|\varepsilon_j)|\varepsilon_i]$ is player i 's conditional expectation of the probability for the opponent player j to choose action a_j given σ_j . A strategy profile $\sigma \in \Sigma$ is a Bayesian Nash equilibrium if, for every player i , $\sigma_i(a_i|\varepsilon_i) > 0$ only if

$$\sum_{a_j \in A_j} g_i(a_i, a_j) E[\sigma_j(a_j|\varepsilon_j)|\varepsilon_i] + \varepsilon_i(a_i) \geq \sum_{a_j \in A_j} g_i(a'_i, a_j) E[\sigma_j(a_j|\varepsilon_j)|\varepsilon_i] + \varepsilon_i(a'_i)$$

for all $a'_i \in A_i$.

Let σ be a Bayesian Nash equilibrium. Suppose that ε_1 and ε_2 are independent. Consider a probability distribution of action profiles, $\mu \in \Delta(A)$, generated by σ : $\mu(a) = E[\sigma_1(a_1|\varepsilon_1) \times \sigma_2(a_2|\varepsilon_2)]$. Mckelvey and Palfrey (1995) called μ a quantal response equilibrium.¹ They showed that, under a specific distributional assumption, if the variance of ε_i is large enough, then a quantal response equilibrium is unique.

¹The definition here is equivalent to the original definition in Mckelvey and Palfrey (1995) which did not use the incomplete information game framework.

Theorem 1 *Suppose that $\varepsilon_i(a_i)$ is independently and identically distributed according to a log Weibul distribution for all $a_i \in A_i$ and $i \in N$: the probability density function of $(\varepsilon_1, \varepsilon_2)$ is $p(\varepsilon_1, \varepsilon_2) = \prod_{i \in N} \prod_{a_i \in A_i} \beta \exp(-\beta \varepsilon_i(a_i)) \exp(-\exp(\beta \varepsilon_i(a_i)))$. If β is small enough, i.e., the variance is large enough, then a quantal response equilibrium is unique.*

In this paper, we assume that private signals are jointly normally distributed with the following symmetric structure: for $i \neq j$,

- every component has mean zero: $E(\varepsilon_i(\alpha)) = E(\varepsilon_i(\beta)) = 0$,
- every component has the same variance: $\text{Var}(\varepsilon_i(\alpha)) = \text{Var}(\varepsilon_i(\beta)) = \rho^2$,
- components of the same player are independent: $\text{Cov}(\varepsilon_i(\alpha), \varepsilon_i(\beta)) = 0$,
- components of the different players are correlated:

$$\text{Cov}(\varepsilon_i(\alpha), \varepsilon_j(\alpha)) = \text{Cov}(\varepsilon_i(\beta), \varepsilon_j(\beta)) = \phi \rho^2, \quad \text{Cov}(\varepsilon_i(\alpha), \varepsilon_j(\beta)) = \psi \rho^2.$$

The parameter ϕ is a correlation coefficient of $\varepsilon_i(\alpha)$ and $\varepsilon_j(\alpha)$ and that of $\varepsilon_i(\beta)$ and $\varepsilon_j(\beta)$. The parameter ψ is a correlation coefficient of $\varepsilon_i(\alpha)$ and $\varepsilon_j(\beta)$. We assume that $|\phi| < 1$, $|\psi| < 1$, $|\phi + \psi| < 1$, and $|\phi - \psi| < 1$ for the multivariate normal distribution to be well defined. The tuple (ρ, ϕ, ψ) parametrizes the information structure of the game.

3 Result

The following theorem provides a sufficient condition for the uniqueness of Bayesian Nash equilibria.

Theorem 2 *If*

$$\gamma(\rho, \phi, \psi) \equiv \frac{1}{\rho} \sqrt{\frac{1 - (\phi - \psi)}{1 + (\phi + \psi)}} < \frac{2\sqrt{\pi}}{x - u + y - v} \quad (1)$$

then every Bayesian Nash equilibrium σ is such that $\sigma_i(\alpha|\varepsilon_i) = 1$ if $\varepsilon_i(\alpha) - \varepsilon_i(\beta) > r^$ and $\sigma_i(\alpha|\varepsilon_i) = 0$ if $\varepsilon_i(\alpha) - \varepsilon_i(\beta) < r^*$ for $i \in \{1, 2\}$ where r^* solves*

$$\Phi\left(\frac{\gamma(\rho, \phi, \psi)r}{\sqrt{2}}\right) = \frac{r + x - u}{x - u + y - v}$$

and Φ is a cumulative probability distribution function of a standard normal distribution.

The theorem asserts that player's behavior is the same in all Bayesian Nash equilibria almost surely, except when $\varepsilon_i(\alpha) - \varepsilon_i(\beta) = r^*$. In this sense, the equilibrium is unique. Note that, if $\gamma(\rho, \phi, \psi) \simeq 0$, then $\Phi\left(\frac{\gamma(\rho, \phi, \psi)r}{\sqrt{2}}\right) \simeq 1/2$, and thus $r^* \simeq (y - v + u - x)/2$.

To understand what (1) implies, consider a simple case with $\psi = 0$, $x = 2$, $y = 1$, $u = v = 0$. Equation (1) is reduced to $\rho^2 > (1 - \phi)/(1 + \phi) \times 9/(4\pi)$. The right-hand side is a decreasing function of ϕ , and its graph is in Figure 1. For the equilibrium to be




Figure 1: $\rho^2 = (1 - \phi)/(1 + \phi) \times 9/(4\pi)$

unique, the variance ρ^2 must be large if the correlation ϕ is low. This corresponds to the uniqueness of quantal response equilibria. Note that if $\rho^2 > 9/4\pi$ then the equilibrium is unique for any $0 \leq \phi < 1$. Also, for the equilibrium to be unique, the correlation ϕ must be high if the variance ρ^2 is small. This corresponds to the uniqueness of equilibria in global games.

To prove the theorem, we use the iterated deletion of dominated strategies, following Carlsson and van Damme (1993). We briefly see the intuition of how it proceeds.

If $\varepsilon_i(\alpha) - \varepsilon_i(\beta) > \max\{u - x, y - v\}$, then player i 's rational choice is α because α dominates β . This is the first round of the iterated deletion. In the remaining rounds, we have to show that player i 's rational choice is α as far as $\varepsilon_i(\alpha) - \varepsilon_i(\beta) > r^*$.

Suppose that $\varepsilon_i(\alpha) - \varepsilon_i(\beta) \simeq \max\{u - x, y - v\}$. Since $\max\{u - x, y - v\} \geq (u - x + y - v)/2$, α risk-dominates β . Thus, if player i believes that the probability of the opponent to choose α is close to or more than $1/2$, then player i 's rational choice is α , and thus we have another round of the iterated deletion. This is true through two effects: the large variance effect and the high correlation effect.

If the variance ρ^2 is very large, player i believes that, with probability close to $1/2$, the difference of the opponent's signal $\varepsilon_j(\alpha) - \varepsilon_j(\beta)$ is positive and very large ($> \max\{u - x, y - v\}$), and α dominates β for the opponent player. This implies that the probability of the opponent to choose α is close to $1/2$.

If the correlation ϕ is very high, player i believes that $\varepsilon_j(\alpha) - \varepsilon_j(\beta)$ is centered around $\varepsilon_i(\alpha) - \varepsilon_i(\beta) \simeq \max\{u - x, y - v\}$, and thus with probability close to $1/2$, $\varepsilon_j(\alpha) - \varepsilon_j(\beta) > \max\{u - x, y - v\}$. This implies that the probability of the opponent to choose α is close to $1/2$.

Repeating the similar argument, we can complete the iterated deletion of dominated strategies and obtain Theorem 2. The formal proof is in the appendix.

4 Discussion

As we have seen in the previous section, if the variance is large enough, then an equilibrium is unique for any non-negative correlation. In contrast, it is not the case in global games. We discuss this issue by considering a global game with parameters similar to our setup.

Suppose that payoffs are randomly determined and each player makes an independent noisy observation. The random payoffs are given by the following table where $\xi(\alpha)$ and $\xi(\beta)$ are i.i.d. normal random variables with mean zero.

	α	β
α	$x + \xi(\alpha), x + \xi(\alpha)$	$v + \xi(\alpha), u + \xi(\beta)$
β	$u + \xi(\beta), v + \xi(\alpha)$	$y + \xi(\beta), y + \xi(\beta)$

Each player knows the constants $x, y, u,$ and $v,$ but does not directly observe his own payoffs, which is different from our setup. Player i 's private signal is an independent noisy observation of the above random payoffs, or equivalently, $\varepsilon_i = (\varepsilon_i(\alpha), \varepsilon_i(\beta)) = (\xi(\alpha) + \eta_i(\alpha), \xi(\beta) + \eta_i(\beta))$ where $\eta_i(\alpha)$ and $\eta_i(\beta)$ are i.i.d. normal random variables with mean zero.

Carlsson and van Damme (1993) called the above setup a global game and obtained a sufficient condition for the uniqueness in terms of the variances of ξ and η_i : if the variance of η_i is small enough compared to that of ξ , then a Bayesian Nash equilibrium is unique. Alternatively, we study the global game in terms of the variance and the correlation of the private signals ε_1 and ε_2 . That is, we parametrize the probability distribution of the private signals in terms of $\rho, \phi > 0$ such that $\rho^2 = \text{Var}(\varepsilon_i(a_i)) = \text{Var}(\xi(a_i)) + \text{Var}(\eta_i(a_i))$ for $a_i \in A_i$ and $\phi = \text{Cov}(\varepsilon_i(a_i), \varepsilon_j(a_j)) / \rho^2 = \text{Var}(\xi(a_i)) / \rho^2$ for $a_i = a_j$ and $i \neq j$. Note that $\text{Cov}(\varepsilon_i(\alpha), \varepsilon_i(\beta)) = \text{Cov}(\varepsilon_i(\alpha), \varepsilon_j(\beta)) = 0$. Given ρ and ϕ , the probability distribution of the private signals is the same as that in our setup with $\psi = 0$. By the standard calculation of multivariate normal distributions, we have $E[\xi(a_i) | \varepsilon_i] = \phi \varepsilon_i(a_i)$ for $a_i \in A_i$, and the interim *expected* payoffs for each player are given by the following table.

	α	β
α	$x + \phi \varepsilon_1(\alpha), x + \phi \varepsilon_2(\alpha)$	$v + \phi \varepsilon_1(\alpha), u + \phi \varepsilon_2(\beta)$
β	$u + \phi \varepsilon_1(\beta), v + \phi \varepsilon_2(\alpha)$	$y + \phi \varepsilon_1(\beta), y + \phi \varepsilon_2(\beta)$

Let us obtain a sufficient condition for the uniqueness in terms of ρ and ϕ . While we can obtain the condition by rewriting the result reported in Morris and Shin (2002), Theorem 2 also leads us to the condition. We take the latter approach, which will help us to understand the driving force of the differences and similarities between our result and the global game analysis. First, observe that the global game has the same set of Bayesian Nash equilibria as that of an incomplete information game such that player i 's private signal is ε_i and the *actual* payoffs for each player are given by the above table. Second, observe that the set of Bayesian Nash equilibria is unchanged even if we multiply all payoffs by a constant $1/\phi$, which induces an incomplete information game such that




Figure 2: $\rho^2 = (1 - \phi)/(\phi^2(1 + \phi)) \times 9/(4\pi)$

player i 's private signal is ε_i and the actual payoffs for each player are given by the following table.

	α	β
α	$x/\phi + \varepsilon_1(\alpha), x/\phi + \varepsilon_2(\alpha)$	$v/\phi + \varepsilon_1(\alpha), u/\phi + \varepsilon_2(\beta)$
β	$u/\phi + \varepsilon_1(\beta), v/\phi + \varepsilon_2(\alpha)$	$y/\phi + \varepsilon_1(\beta), y/\phi + \varepsilon_2(\beta)$

Third, apply Theorem 2 to this game and translate it for the global game taking account of the equivalence of the equilibrium sets. Then, we conclude that the global game has a unique Bayesian Nash equilibrium if

$$\frac{1}{\phi\rho} \sqrt{\frac{1-\phi}{1+\phi}} < \frac{2\sqrt{\pi}}{x-u+y-v}. \quad (2)$$

To understand what (2) implies, consider a simple case with $x = 2$, $y = 1$, $u = v = 0$ as before. Equation (2) is reduced to $\rho^2 > (1 - \phi)/(\phi^2(1 + \phi)) \times 9/(4\pi)$, and its graph is in Figure 2. Compare Figure 1 and Figure 2. They are similar in the sense that the infimum of the variance necessary for the uniqueness is decreasing in ϕ . But there is an obvious distinction. The infimum in the global game is larger because $(1 - \phi)/(\phi^2(1 + \phi)) > (1 - \phi)/(1 + \phi)$, and it goes to infinity as ϕ goes to zero. This implies that the very low correlation leads to the failure of uniqueness even if the variance is very large.

Appendix

Proof of Theorem 2. For $i \neq j$, let $F_i(\sigma_j, \varepsilon_i)$ be the difference of an expected payoff of player i with ε_i taking α and that of the same player taking β when the other player j 's strategy is $\sigma_j \in \Sigma_j$:

$$F_i(\sigma_j, \varepsilon_i) = \sum_{a_j \in A_j} g_i(\alpha, a_j) E[\sigma_j(a_j | \varepsilon_j) | \varepsilon_i] + \varepsilon_i(\alpha) - \sum_{a_j \in A_j} g_i(\beta, a_j) E[\sigma_j(a_j | \varepsilon_j) | \varepsilon_i] - \varepsilon_i(\beta).$$

Note that if $\sigma_j(\alpha|\varepsilon_j) \geq \sigma'_j(\alpha|\varepsilon_j)$ for all ε_j then $F_i(\sigma_j, \varepsilon_i) \geq F_i(\sigma'_j, \varepsilon_i)$ for all ε_i .

Let $\delta_j^r \in \Sigma_j$ be a switching strategy such that $\delta_j^r(\alpha|\varepsilon_j) = 1$ if $\varepsilon_j(\alpha) - \varepsilon_j(\beta) \geq r$ and $\delta_j^r(\alpha|\varepsilon_j) = 0$ if $\varepsilon_j(\alpha) - \varepsilon_j(\beta) < r$. Then,

$$F_i(\delta_j^r, \varepsilon_i) = -(x - u + y - v) \Pr[\varepsilon_j(\alpha) - \varepsilon_j(\beta) < r | \varepsilon_i] + x - u + \varepsilon_i(\alpha) - \varepsilon_i(\beta).$$

By the standard calculation of multivariate normal distributions, we obtain

$$\Pr[\varepsilon_j(\alpha) - \varepsilon_j(\beta) < r | \varepsilon_i] = \Phi\left(\frac{r - (\phi - \psi)(\varepsilon_i(\alpha) - \varepsilon_i(\beta))}{\rho\sqrt{2(1 - (\phi - \psi)^2)}}\right)$$

where Φ is a cumulative probability distribution function of a standard normal distribution. Writing $r_i = \varepsilon_i(\alpha) - \varepsilon_i(\beta)$, we have

$$F_i(\delta_j^r, \varepsilon_i) = -(x - u + y - v)\Phi\left(\frac{r - (\phi - \psi)r_i}{\rho\sqrt{2(1 - (\phi - \psi)^2)}}\right) + x - u + r_i.$$

Let $G(r, r_i) = F_i(\delta_j^r, \varepsilon_i)$. Then, $G(r, r_i)$ is strictly decreasing in r , strictly increasing in r_i , and $G(r, r)$ is strictly increasing in r because

$$\begin{aligned} \frac{\partial}{\partial r} G(r, r_i) &= -\frac{x - u + y - v}{\rho\sqrt{2(1 - (\phi - \psi)^2)}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{r - (\phi - \psi)r_i}{\rho\sqrt{2(1 - (\phi - \psi)^2)}}\right)^2\right) < 0, \\ \frac{\partial}{\partial r_i} G(r, r_i) &= \frac{(x - u + y - v)(\phi - \psi)}{\rho\sqrt{2(1 - (\phi - \psi)^2)}} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{r - (\phi - \psi)r_i}{\rho\sqrt{2(1 - (\phi - \psi)^2)}}\right)^2\right) + 1 > 0, \\ \frac{d}{dr} G(r, r) &= \frac{\partial}{\partial r} G(r, r_i) \Big|_{r_i=r} + \frac{\partial}{\partial r_i} G(r, r_i) \Big|_{r_i=r} \\ &= -\left(\frac{(x - u + y - v)(1 - (\phi - \psi))}{2\sqrt{\pi}\rho\sqrt{1 - (\phi - \psi)^2}}\right) \exp\left(-\frac{1}{2}\left(\frac{r - (\phi - \psi)r}{\rho\sqrt{2(1 - (\phi - \psi)^2)}}\right)^2\right) + 1 \\ &\geq -\frac{(x - u + y - v)}{2\sqrt{\pi}\rho} \sqrt{\frac{1 - (\phi - \psi)}{1 + (\phi + \psi)}} + 1 > 0 \end{aligned}$$

where the last inequality is due to the assumption of (1). Let r^* be such that $G(r^*, r^*) = 0$. Note that r^* is unique and it is the same as r^* in the statement of the theorem.

For $s \in \mathbb{R}$, consider a subset of Σ_i such that

$$\begin{aligned} \bar{\Sigma}_i(s) &= \{\sigma_i \in \Sigma_i \mid \sigma_i(\alpha|\varepsilon_i) = 1 \text{ if } \varepsilon_i(\alpha) - \varepsilon_i(\beta) > s\}, \\ \underline{\Sigma}_i(s) &= \{\sigma_i \in \Sigma_i \mid \sigma_i(\alpha|\varepsilon_i) = 0 \text{ if } \varepsilon_i(\alpha) - \varepsilon_i(\beta) < s\}. \end{aligned}$$

Let $\bar{s}^1 = y - v$. If $\varepsilon_i(\alpha) - \varepsilon_i(\beta) > \bar{s}^1$ then $F_i(\sigma_j, \varepsilon_i) > 0$ for all $\sigma_j \in \Sigma_j$. This implies that any $\sigma_i \notin \bar{\Sigma}_i(\bar{s}^1)$ is strictly dominated by some $\sigma'_1 \in \bar{\Sigma}_i(\bar{s}^1)$. Let $\underline{s}^1 = -(x - u)$. If $\varepsilon_i(\alpha) - \varepsilon_i(\beta) < \underline{s}^1$ then $F_i(\sigma_j, \varepsilon_i) < 0$ for all $\sigma_j \in \Sigma_j$. This implies that any $\sigma_i \notin \underline{\Sigma}_i(\underline{s}^1)$

is strictly dominated by some $\sigma'_i \in \underline{\Sigma}_i(\underline{s}^1)$. By the first round of deletion, we obtain $\Sigma_i^1 = \overline{\Sigma}_i(\overline{s}^1) \cap \underline{\Sigma}_i(\underline{s}^1)$ for $i \in N$.

Let \overline{s}^2 be such that $G(\overline{s}^1, \overline{s}^2) = 0$. Note that $\overline{s}^2 < \overline{s}^1$ because $G(\overline{s}^1, \overline{s}^1) > 0$. Thus, if $\varepsilon_i(\alpha) - \varepsilon_i(\beta) > \overline{s}^2$ and $\sigma_j \in \Sigma_j^1 \subseteq \overline{\Sigma}_j(\overline{s}^1)$ then $F_i(\sigma_j, \varepsilon_i) \geq F_i(\delta_j^{\overline{s}^1}, \varepsilon_i) = G(\overline{s}^1, \varepsilon_i(\alpha) - \varepsilon_i(\beta)) > G(\overline{s}^1, \overline{s}^2) = 0$. Thus, any $\sigma_i \notin \overline{\Sigma}_i(\overline{s}^2)$ is strictly dominated by some $\sigma'_i \in \overline{\Sigma}_i(\overline{s}^2)$ against any $\sigma_j \in \Sigma_j^1$. Symmetrically, let \underline{s}^2 be such that $G(\underline{s}^1, \underline{s}^2) = 0$. Note that $\underline{s}^2 > \underline{s}^1$ because $G(\underline{s}^1, \underline{s}^1) < 0$. Thus, if $\varepsilon_i(\alpha) - \varepsilon_i(\beta) < \underline{s}^2$ and $\sigma_j \in \Sigma_j^1 \subseteq \underline{\Sigma}_j(\underline{s}^1)$ then $F_i(\sigma_j, \varepsilon_i) \leq F_i(\delta_j^{\underline{s}^1}, \varepsilon_i) = G(\underline{s}^1, \varepsilon_i(\alpha) - \varepsilon_i(\beta)) < G(\underline{s}^1, \underline{s}^2) = 0$. Thus, any $\sigma_i \notin \underline{\Sigma}_i(\underline{s}^2)$ is strictly dominated by some $\sigma'_i \in \underline{\Sigma}_i(\underline{s}^2)$ against any $\sigma_j \in \Sigma_j^1$. By the second round of deletion, we obtain $\Sigma_i^2 = \overline{\Sigma}_i(\overline{s}^2) \cap \underline{\Sigma}_i(\underline{s}^2)$ for $i \in N$.

We repeat this procedure: construct a decreasing sequence $\{\overline{s}^k\}_{k=1}^\infty$ and an increasing sequence $\{\underline{s}^k\}_{k=1}^\infty$ such that $G(\overline{s}^k, \overline{s}^{k+1}) = G(\underline{s}^k, \underline{s}^{k+1}) = 0$, and define $\Sigma_i^k = \overline{\Sigma}_i(\overline{s}^k) \cap \underline{\Sigma}_i(\underline{s}^k)$. We obtain Σ_i^k by the k -th round of deletion. Since $\lim_{k \rightarrow \infty} \overline{s}^k = \lim_{k \rightarrow \infty} \underline{s}^k = r^*$, we have

$$\begin{aligned} \Sigma_i^* &\equiv \bigcap_{k=1}^{\infty} \Sigma_i^k = \overline{\Sigma}_i(r^*) \cap \underline{\Sigma}_i(r^*) \\ &= \{\sigma_i \in \Sigma_i \mid \sigma_i(\alpha|\varepsilon_i) = 1 \text{ if } \varepsilon_i(\alpha) - \varepsilon_i(\beta) > r^*, \sigma_i(\alpha|\varepsilon_i) = 0 \text{ if } \varepsilon_i(\alpha) - \varepsilon_i(\beta) < r^*\}. \end{aligned}$$

Any Bayesian Nash equilibrium must be in $\Sigma^* = \Sigma_1^* \times \Sigma_2^*$, which completes the proof. ■

References

- Carlsson, H., and van Damme, E. (1993). “Global Games and Equilibrium Selection,” *Econometrica* **61**, 989–1018.
- Harsanyi, J. C. (1973). “Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points,” *Int. J. Game Theory* **2**, 1–23.
- McKelvey, D., and Palfrey, T. R. (1995). “Quantal Response Equilibria for Normal Form Games,” *Games Econ. Behav.* **10**, 6–38.
- Morris, S., and Shin, H. S. (2002). “Global Games: Theory and Applications,” in *Advances in Economics and Econometrics: Theory and Applications, Proceedings of the Eighth World Congress of the Econometric Society* (M. Dewatripont, L. Hansen, and S. Turnovsky Eds.), pp. 56–114. Cambridge Univ. Press.
- Morris, S., and Shin, H. S. (2005). “Heterogeneity and Uniqueness in Interaction Games,” in *The Economy as an Evolving Complex System III* (L. Blume, and S. Durlauf Eds.), pp. 207–242. Oxford Univ. Press.