

An incomplete information justification of symmetric equilibrium in symmetric games

Christoph Kuzmics*, and Brian Rogers†

November 19, 2010

Abstract

Consider a symmetric 2-player game of complete information. Consider an arbitrary Bayesian extension of that game with payoff-irrelevant types, independent random matching, and anonymity (private types). We show that, in this setting, while strategies in a Bayesian Nash equilibrium of that game can differ across types, aggregate play in any such equilibrium must coincide with a symmetric Nash equilibrium of the complete information game. This justifies the interpretation of certain data, including many laboratory experiments, as arising from a symmetric equilibrium, even when asymmetric equilibria exist and, in addition, subjects may be heterogeneous.

JEL codes: C72, C90

Keywords: symmetry, Bayesian game, laboratory experiments

I. Introduction

Consider an experimentalist whose subjects are playing a symmetric 2-player game in the lab. The analyst is interested in what behavior can arise at equilibrium. Can the

*Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, c-kuzmics@kellogg.northwestern.edu.

†Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, b-rogers@kellogg.northwestern.edu.

data be viewed as arising not from an arbitrary equilibrium but, more precisely, from a symmetric equilibrium?¹ The analyst allows for the possibility that different types of players, though all rational, may approach the game differently. There may be, for instance, men and women among the players in the lab. There may be men and women who consider themselves tough, while some consider themselves weak, etc. Is it possible that these different types of players end up playing different strategies in equilibrium?

Consider generally, a symmetric game, being played by two players in the field. Under what conditions can we guarantee that their behavior is as if they were playing a symmetric equilibrium? Under what conditions is this not necessarily the case? The answers we provide to these questions are already understood, at least informally, to a certain extent. The value of our formalization, while not technically demanding, is to isolate the conditions that imply that behavior will conform to a symmetric equilibrium, even when there are asymmetric equilibria and, moreover, players may be heterogeneous.

In order to address these questions we consider a general symmetric 2-player game, which we call the *base game*, and enrich it by adding a set of, and distribution over, types of players, yielding a *Bayesian extension* of the base game. Call types *payoff irrelevant* if every type shares the same payoffs as given by the base game. The Bayesian extension is said to satisfy *independent random matching* if the types of the two players are chosen independently. It is said to satisfy *anonymity* if players' types are private information.

We show that under the assumptions of payoff-irrelevant types, independent random matching, and anonymity, every Bayesian equilibrium of every Bayesian extension of the base game induces aggregate play (as observed by the analyst) to be as in a symmetric Nash equilibrium of the base game. It is not necessarily true, however, that every type of player uses the same strategy. In equilibrium, however, all strategies that are used must yield the same payoff.

Anonymity and independent random matching are satisfied in many experimental designs. In particular, the conventions of assigning subjects into pairs uniformly at random, and preventing subjects from observing the identity of their opponent imply that these conditions are (approximately) met.

The third assumption that all types share the same payoff is not as easy to ensure, as individuals may have heterogeneous (non-linear) utility in money, or care about outcomes beyond their own payoff. If player types have heterogeneous preferences which are close to the preferences in the base game, then every Bayesian equilibrium of any Bayesian

¹In many cases such a conclusion would reduce or even eliminate the multiplicity problem.

extension which satisfies independent random matching and anonymity must be close to a symmetric Nash equilibrium of the base game. We also discuss the effects of relaxing the two assumptions of independent random matching and anonymity.

II. Model and Main Result

Let $\Gamma = (S, u)$ be a symmetric 2-player game with finite *action space* S and payoff function $u : S \times S \rightarrow \mathbb{R}$ (or $u : S^n \rightarrow \mathbb{R}$) with the interpretation that $u(s, s')$ is the payoff a player gets when playing s against an opponent playing s' . This game Γ shall be called the *base game*. Let $\Delta(S)$ denote the set of mixed actions. Let u be extended to mixed actions in the usual expected utility way.

We shall now provide a definition of an incomplete information game that is built from the base game. Let Θ be a finite set of *types* with typical element θ . Let $\mu^* \in \Delta(\Theta)$ denote the probability distribution over types. Let $\mu \in \Delta(\Theta \times \Theta)$ denote the *matching technology*, i.e., the probability distribution that describes the frequency with which a pair of player types θ, θ' is drawn. We assume that the marginal of μ in both arguments is μ^* , so that types are drawn in proportion to their frequency in the population. Let K be a finite set of *signals* with typical element κ . For any type $\theta \in \Theta$ let $\nu_\theta \in \Delta(K)$ denote *conditional signal distribution* for type θ . That is, $\nu_\theta(\kappa)$ is the probability that a type θ generates signal κ . Signals are perfectly observable.² Given this framework, we define a *strategy* by $\sigma : \Theta \times K \rightarrow \Delta(S)$ with the interpretation that $\sigma_{\theta, \kappa}$ is the (mixed) action a player of type θ employs when encountering an opponent who so happened to generate signal κ . This completes the description of a *Bayesian extension*, $\Gamma^B = [\Gamma = (S, u), (\Theta, \mu), (K, \nu)]$, of the base game Γ .

A Bayesian extension, Γ^B , satisfies *independent random matching* if $\mu(\theta, \theta') = \mu^*(\theta)\mu^*(\theta')$ for all $\theta, \theta' \in \Theta$. That is, the matching technology is simply the product of its (equal) marginals. The Bayesian extension satisfies *anonymity* if $\nu_\theta(\kappa) = \nu_{\theta'}(\kappa)$ for all $\theta, \theta' \in \Theta$ and all $\kappa \in K$. Alternatively, in such a case we can simply write K is a singleton. The important aspect is that no type of player can infer anything about her opponent's type from the observed signal.

We are interested in the following quantities derived from the strategies used in a Bayesian extension of some base game. Let $y_\sigma(s, s')$ denote the aggregate probability that a randomly matched pair of players play action profile (s, s') . Thus, $y_\sigma \in \Delta(S \times S)$. We shall call y_σ the *aggregate distribution of play*. Let y_σ^* denote its marginal (which must be

²There is little loss of generality in this, as we shall see momentarily.

equal in both arguments). Furthermore, let $z_{\sigma,\theta,\kappa}(s)$ denote the probability that a player of type θ who observes signal κ faces action s . Thus, $z_{\sigma,\theta,\kappa} \in \Delta(S)$. That is, $z_{\sigma,\theta,\kappa}(s)$ is the probability that a random (under μ and ν) opponent of this player type plays action s .

We now provide an immediate, but crucial, lemma.

LEMMA 1: *Let the Bayesian extension, $\Gamma^B = [\Gamma = (S, u), (\Theta, \mu), (K, \nu)]$ satisfy independent random matching and anonymity. Fix an arbitrary strategy σ . Then*

1. *every type faces the same distribution of play, i.e. $z_{\sigma,\theta,\kappa} = z_\sigma$ for all $\theta \in \Theta$ and all $\kappa \in K$,*
2. *the aggregate distribution of play is the product of its two (equal) marginals, i.e. $y_\sigma(s, s') = y_\sigma^*(s)y_\sigma^*(s')$ for all $s, s' \in S$,*
3. *and the above two are identical, i.e. $z_\sigma = y_\sigma^*$.*

Proof: See Appendix.

Roughly, the idea of Lemma 1 is that, since matching is independent and signals are uninformative, each player type, when taking her action, faces a distribution of opponent types that is given by the population frequencies μ^* . Thus, each player type faces the same distribution of play, which then must be equal to the marginal distribution over actions, y_σ^* . Since under our assumptions the actions of players must be statistically independent, the joint distribution over actions, y_σ , is simply the product of its marginals.

Notice that the lemma relies in no way on payoffs or incentives. It is a statement about what kind of play is (not) technologically feasible given the assumptions on the matching process and type observability. With these observations in hand, however, it is easy to deliver the main result that aggregate equilibrium play in a Bayesian extension must conform to that of a symmetric equilibrium of the base game.

A *Nash equilibrium* of the base game is a pair of mixed actions $(x, x') \in \Delta(S)^2$ such that $u(x, x') \geq u(s, x')$ and $u(x', x) \geq u(s, x)$ for all $s \in S$. It is a *symmetric Nash equilibrium* if, in addition, $x = x'$. We shall then also call $x \in \Delta(S)$ a symmetric Nash equilibrium (action) of the base game. Every symmetric game has a symmetric Nash equilibrium.³

In order to talk about equilibria in a Bayesian extension of the base game, we need to introduce payoffs to all types. For the main result we shall assume that types are

³This follows from Theorem 2 in Nash (1951).

payoff-irrelevant, which is to say that every type of every player shares the same payoff function, which is the one given by the base game. Generally, though, we introduce the function $U : \Theta \times S \times S \rightarrow \mathbb{R}$, describing type-dependent payoffs. Utilities are extended to mixed strategies in the usual way. Let, thus, $U_\theta(s, s')$ denote the payoff of player type θ if this player plays action s while her opponent uses action s' . Given a strategy profile $\sigma : \Theta \times K \rightarrow \Delta(S)$ let $U_\theta(\sigma_{\theta,\kappa}, z_{\sigma,\theta,\kappa})$ denote the expected payoff to a player of type θ upon observing the signal κ under strategy profile σ .

A *Bayesian Nash equilibrium* profile σ for the Bayesian extension Γ^B is such that $U_\theta(\sigma_{\theta,\kappa}, z_{\sigma,\theta,\kappa}) \geq U_\theta(s, z_{\sigma,\theta,\kappa})$ for all $s \in S$, all $\theta \in \Theta$, and all $\kappa \in K$. Again, a Bayesian Nash equilibrium always exists.

THEOREM 1: *Let $\Gamma = (S, u)$ be a symmetric 2-player base game. Then for any Bayesian extension with payoff irrelevant types, independent random matching, and anonymity, and any Bayesian Nash equilibrium σ of that game we have that all types receive the same expected payoff and the induced aggregate mixed action y_σ^* is a symmetric Nash equilibrium of the base game.*

Proof: By Lemma 1 we have that $z_{\sigma,\theta,\kappa} = z_\sigma = y_\sigma^*$ for all $\theta \in \Theta$ and all $\kappa \in K$. Thus, any player of any type faces the same distribution of opponent play given by y_σ^* . Thus, every type in the equilibrium σ must be playing a best response to y_σ^* . Thus, all types receive the same expected payoff. Furthermore as y_σ^* is simply a convex combination of best responses to itself, it must be a best response to itself. Thus, y_σ^* is a symmetric Nash equilibrium.⁴ QED

III. Discussion

In this section we discuss Theorem 1 and the effect of relaxing the assumptions of payoff-irrelevant types, independent random matching, and anonymity, by means of the simplest 2-player symmetric game of interest, which is given by

	H	L	
H	0	2	.
L	1	0	

This game can be viewed as a battle-of-the-sexes game, a chicken game, or a Hawk-Dove game. Verbally the game is described as follows. Every player has two actions, H (for

⁴Notice that, while the notation would become more cumbersome, the argument of (Lemma 1 and) Theorem 1 extends directly to the case of n -player symmetric normal form games.

high), and L (for low). Let us call a typical player “You”. If you choose H while your opponent chooses L you get a high payoff of 2. If you choose L while your opponent chooses H you get a low payoff of 1. In both other cases you get a payoff of zero. The description of the game for your opponent is exactly the same. Note that this game has a unique symmetric equilibrium, in which you play H with probability $\frac{2}{3}$ and L with probability $\frac{1}{3}$. The game also has two asymmetric Nash equilibria. One in which you play H while your opponent plays L, and another in which you play L while your opponent plays H.

Suppose we enrich this game by considering the following Bayesian extension, in which there are two types, call them Aristocrats and Commoners, which are payoff irrelevant. Suppose also that we have independent random matching and anonymity, so that the conditions of Theorem 1 are satisfied. Suppose further that the marginal distribution over types is such that the frequency of Aristocrats is $\frac{2}{3}$ while the frequency of Commoners is $\frac{1}{3}$. Note that one of the Bayesian equilibria of this Bayesian extension is such that Aristocrats play H while Commoners play L. Given the distribution over types, independent random matching, and anonymity, we have that every player faces the same distribution of actions, in this case $\frac{2}{3}$ on H and $\frac{1}{3}$ on L whether they are Aristocrats or Commoners. Thus, it so happens that the presumed behavior of the two types is indeed in equilibrium. Note that these two types could also randomize very differently in another Bayesian equilibrium. For instance, they could of course both play the symmetric mixed Nash equilibrium. There is in fact a continuum of equilibria with different behavior for the two types.

Now, can we think of Aristocrats as those types of players that always play H (in all Bayesian extensions with independent random matching and anonymity)? The answer is no. Suppose that the frequency of Aristocrats is strictly above $\frac{2}{3}$. Then there is no equilibrium in which Aristocrats play H with probability 1. Thus, types cannot behave equally in all Bayesian extensions. In other words, while every mixed action whose support coincides with that of a symmetric Nash equilibrium, can be played by some type in some Bayesian extension, the only mixed action that can be played by some type in every Bayesian extension is the symmetric Nash equilibrium of the base game.

A. *Heterogeneous preferences*

In the lab the assumptions of independent random matching and anonymity are typically satisfied and, in any case, can certainly be imposed. The third assumption used in Theorem 1, the assumption of payoff-irrelevant types, is not as easy to impose. For instance,

some types of players might have preferences that are not simply linear in money. It is easy to see, however, that as long as all types have preferences that are close to the base game preferences, then Lemma 1 is still exactly true and an approximate version of Theorem 1 still holds: aggregate behavior in this Bayesian extension must be close to a symmetric Nash equilibrium.⁵ If there is heterogeneity in preferences such that preferences are sufficiently diverse a purification result⁶, as in Harsanyi (1973), must hold.⁷ Not only is the aggregate behavior in every Bayesian Nash equilibrium close to a symmetric Nash equilibrium of the base game, but in all of these equilibria almost every type of player uses a pure action.

B. *Type observability*

What happens when anonymity is relaxed? The strongest form of departure from anonymity is perfect observability. That is, when a pair of players are matched, their type profile is common knowledge. In this case, an equilibrium of the Bayesian extension can correspond to any (possibly asymmetric) equilibrium of the base game. A striking example is provided in Armstrong and Duffy (2010), who study the Battle of the Sexes in the online game “World of Warcraft”. In that setting, each player has a public type which is an integer score, which can be thought of as the player’s status in the World of Warcraft. Whenever two players are matched, it is then common knowledge which player has the higher type and which player has the lower type.⁸ One equilibrium of the Bayesian extension, and the one that is empirically observed most frequently, is the asymmetric equilibrium where play is coordinated on the preferred outcome of the higher type player. Note that, in this equilibrium, the aggregate distribution of play is, of course, not a product distribution, as it puts probability one-half on the play (H,L) and another half on (L,H). Note also that its marginal, y^* , (the same in both arguments) attaches probability one-half on H and one-half on L, which does not represent a Nash equilibrium of the base game. The distribution of play different types face is not the same either. When a high status type faces a low status type, the high type faces a z which attaches probability one on L, while

⁵Suppose aggregate play was far away from a symmetric equilibrium. Then, given that differences in preferences are slight, all types of players must have the same (pure) best response, which in itself is not an equilibrium. Thus, the equilibrium must be close to a symmetric Nash equilibrium.

⁶Strictly speaking this requires a continuum of types.

⁷Harsanyi (1973) does not specifically talk about symmetric games. Harsanyi’s (1973) Bayesian extension of a symmetric game would be such that players have roles (row and column) as well as types.

⁸For the sake of this discussion we assume that it is essentially impossible that two players with the same score are matched.

the low type faces a z which attaches probability one on H.

Suppose now that anonymity is only slightly violated. We claim that aggregate play in any Bayesian equilibrium must be close to a symmetric Nash equilibrium of the base game. Note that there are many forms this slight violation of anonymity could take. Let us here focus, for the sake of simplicity, on a variant of the “almost no observability” model employed by Dekel, Ely, and Yilankaya (2007) in the context of the evolution of preferences. Suppose $K = \Theta \cup \{\varphi\}$ and $\nu_\theta(\theta) = \epsilon$, $\nu_\theta(\varphi) = 1 - \epsilon$, and $\nu_\theta(\theta') = 0$ for all $\theta' \neq \theta$. This specification of our general model implies that any player of any type has an ϵ probability of observing the opponent type and a $(1 - \epsilon)$ probability of receiving an uninformative signal φ .⁹ If ϵ is small then the aggregate distribution of play is essentially determined by the behavior of types when facing the uninformative signal. Furthermore, the behavior a player expects from an opponent when facing the uninformative signal, again as ϵ is small, is essentially determined by the behavior of players when they, in turn, face an uninformative signal.

We shall now argue that aggregate play in this specification of the Bayesian extension must be close to the symmetric Nash equilibrium of the base game. That is it must be close to $\frac{2}{3}H$ and $\frac{1}{3}L$ in the game at hand. Suppose σ is a Bayesian equilibrium and aggregate play y_σ^* is not close to the symmetric Nash equilibrium of the base game. Then aggregate type behavior against the uninformative signal is not close to the symmetric Nash equilibrium. Then the distribution of actions faced by a player of any type when observing the uninformative signal is not close to the symmetric Nash equilibrium. This is true because ϵ is small. But then all such players have the same best response, either all H if aggregate behavior favors L, or all L if aggregate behavior favors H. But then aggregating this behavior does not lead to the supposed distribution of play and, thus, σ cannot be a Bayesian equilibrium.

C. *Correlated matching*

Independent random matching in the lab is ensured if subjects are randomly drawn from a large pool such that the exact sample of subjects is not known to the subjects. Independent random matching can actually be slightly violated in the lab if all subjects know the composition of the small chosen group of subjects. For instance, suppose that there are 10 subjects and 5 of them are women, while the other 5 are men. Then a woman does

⁹Calling this signal φ uninformative is justified once one realizes that every player of every type, upon observing this signal, faces the same conditional type distribution, which is simply μ^* .

not face the same distribution of men and woman that a man faces. A woman is, in fact, slightly more likely to face a man, than a man is ($\frac{5}{9}$ versus $\frac{4}{9}$). Could this be a concern?

Generally, what is the impact of relaxing independence of the matching technology. Assume first, for simplicity, that there are two types, $\Theta = \{\theta_1, \theta_2\}$, and that matching is perfectly negatively correlated. Then every (asymmetric) equilibrium of the base game has a corresponding equilibrium in the Bayesian extension. For instance, it is an equilibrium for $\sigma(\theta_1) = H$ and $\sigma(\theta_2) = L$. Cooper, DeJong, Forsythe, and Ross (1993) makes precisely this point. They study an extensive form version of battle-of-the-sexes, in which one player moves first and, without seeing the first player's move, the other player moves second. This game is strategically equivalent to the normal form representation of battle-of-the-sexes given above. But empirically, as they demonstrate, play is very different. Why? In the normal form of the game, as we have shown, play must conform to the unique symmetric equilibrium. But in the extensive form, it is commonly known that one player moves first. The observed convention is that this player chooses his most preferred action and, rationally inferring this, the second mover concedes. Thus, the extensive form representation has exactly the type structure given above with θ_1 being the first mover and θ_2 being the second mover and the matching is perfectly negatively correlated, such that a first mover will always know with certainty that she is facing a second mover and vice versa. The resulting distribution of aggregate play (as well as the action distributions different types face) is (are) exactly the same as in the example of perfect observability above.

Consider a finite type space Θ for the battle of the sexes game above. Suppose anonymity is satisfied. Suppose we do not have independent random matching but have the following matching technology. For all $\theta \in \Theta$ we have $\mu(\theta, \theta) > \mu^*(\theta)^2$. That is, every type of player is more likely to meet their own type than would be the case under the marginal distribution. Suppose also that for all $\theta, \theta' \in \Theta$ with $\theta \neq \theta'$ we have $\mu(\theta, \theta')$ is proportional to $\mu^*(\theta)\mu^*(\theta')$.

In this case we can show that there is a unique Bayesian Nash equilibrium for this Bayesian extension, which is such that all types mix $\frac{2}{3}$ on H and $\frac{1}{3}$ on L. That is, not only is the aggregate play as in the symmetric Nash equilibrium, but every type plays as in the symmetric Nash equilibrium of the base game. The reason is as follows. Suppose two types θ, θ' use (mixed) action $\sigma_\theta \neq \sigma_{\theta'}$. Suppose first, that they are both totally mixing. Then both must be indifferent between H and L. However, note that because of the matching technology, and the fact that both use different mixed actions, both face different action distributions, i.e. $z_{\sigma, \theta} \neq z_{\sigma, \theta'}$. But there is only one action distribution that

makes a player indifferent between H and L. Thus, it cannot be that both are indifferent between H and L.

Now consider the second case in which one type uses a pure action, say, w.l.o.g. type θ plays action H. Suppose the other type plays any mixed action that attaches positive probability to L. But then type θ faces an action distribution that attaches more probability on H than type θ' does, due to the positive correlation in the matching technology. But then if type θ' is indifferent between H and L, type θ must prefer L. Again, we arrive at a contradiction.

Third and finally, suppose that the two types both use the same pure action, say, H. But this could only be a best reply if there are some types who play L with sufficiently high probability. By the previous argument, there cannot be such a type. Thus, we have that all types must use the symmetric mixed Nash equilibrium action of attaching $\frac{2}{3}$ on H and the remaining probability on L.

As our final point, suppose that the matching technology has a small degree of negative correlation as in the example of the 5 men and 5 women in the lab given above. Thus, we have $\mu(\theta, \theta) < \mu^*(\theta)^2$ for all $\theta \in \Theta$. Again, suppose that for all $\theta, \theta' \in \Theta$ with $\theta \neq \theta'$ we have $\mu(\theta, \theta')$ is proportional to $\mu^*(\theta)\mu^*(\theta')$. Then, for exactly the same reason as in the positive correlation case, no two types can be fully, yet differently, mixing between H and L. However, it can now be that some types are mixing while other types play pure actions. In fact, the discussion above shows that, for our example and fixing anonymity, there exists a continuum of equilibria essentially only in the Bayesian extension with independent random matching. Any degree of positive correlation reduces the number of equilibria to the unique symmetric Nash equilibrium of the base game, while any degree of negative correlation reduces the number of equilibria to three.

IV. Appendix: Proof of Lemma 1

Let $P(\theta'|\theta, \kappa)$ denote the probability that a player of type θ , upon observing signal κ , assigns to the event that he faces a player of type θ' . Given anonymity we can write $\nu_\theta(\kappa) = \nu(\kappa)$ for all $\theta \in \Theta$. Given also independent random matching, we have that

$P(\theta'|\theta, \kappa) = \mu^*(\theta')$ for all $\theta, \theta' \in \Theta$ and all $\kappa \in K$. Thus,

$$\begin{aligned} z_{\sigma, \theta, \kappa} &= \sum_{\hat{\theta} \in \Theta} \sum_{\hat{\kappa} \in K} \sigma(\hat{\theta}, \hat{\kappa}) P(\hat{\theta}|\theta, \kappa) \nu_{\theta}(\hat{\kappa}) \\ &= \sum_{\hat{\theta} \in \Theta} \sum_{\hat{\kappa} \in K} \sigma(\hat{\theta}, \hat{\kappa}) \mu^*(\hat{\theta}) \nu(\hat{\kappa}) \\ &\equiv z_{\sigma}, \end{aligned}$$

proving part 1.

Next, the marginal distribution over actions is given by

$$\begin{aligned} y_{\sigma}^* &= \sum_{\hat{\theta}} \sum_{\hat{\theta}'} \sum_{\hat{\kappa}} \mu(\hat{\theta}, \hat{\theta}') \nu_{\hat{\theta}'}(\hat{\kappa}) \sigma(\hat{\theta}, \hat{\kappa}) \\ &= \sum_{\hat{\theta}} \sum_{\hat{\theta}'} \sum_{\hat{\kappa}} \mu^*(\hat{\theta}) \mu^*(\hat{\theta}') \nu(\hat{\kappa}) \sigma(\hat{\theta}, \hat{\kappa}) \\ &= \sum_{\hat{\theta}} \sum_{\hat{\kappa}} \mu^*(\hat{\theta}) \nu(\hat{\kappa}) \sigma(\hat{\theta}, \hat{\kappa}) \sum_{\hat{\theta}'} \mu^*(\hat{\theta}') \\ &= \sum_{\hat{\theta}} \sum_{\hat{\kappa}} \mu^*(\hat{\theta}) \nu(\hat{\kappa}) \sigma(\hat{\theta}, \hat{\kappa}), \end{aligned}$$

showing that $z_{\sigma} = y_{\sigma}^*$.

Finally, the joint distribution over outcomes is

$$\begin{aligned} y_{\sigma}(s, s') &= \sum_{\hat{\theta}} \sum_{\hat{\theta}'} \sum_{\hat{\kappa}} \sum_{\hat{\kappa}'} \mu(\hat{\theta}, \hat{\theta}') \nu_{\hat{\theta}'}(\hat{\kappa}) \nu_{\hat{\theta}}(\hat{\kappa}') \sigma(\hat{\theta}, \hat{\kappa})(s) \sigma(\hat{\theta}', \hat{\kappa}')(s') \\ &= \sum_{\hat{\theta}} \sum_{\hat{\theta}'} \sum_{\hat{\kappa}} \sum_{\hat{\kappa}'} \mu^*(\hat{\theta}) \mu^*(\hat{\theta}') \nu(\hat{\kappa}) \nu(\hat{\kappa}') \sigma(\hat{\theta}, \hat{\kappa})(s) \sigma(\hat{\theta}', \hat{\kappa}')(s') \\ &= \left(\sum_{\hat{\theta}} \sum_{\hat{\kappa}} \mu^*(\hat{\theta}) \nu(\hat{\kappa}) \sigma(\hat{\theta}, \hat{\kappa})(s) \right) \left(\sum_{\hat{\theta}'} \sum_{\hat{\kappa}'} \mu^*(\hat{\theta}') \nu(\hat{\kappa}') \sigma(\hat{\theta}', \hat{\kappa}')(s') \right) \\ &= y_{\sigma}^*(s) y_{\sigma}^*(s'), \end{aligned}$$

completing the proof. QED

References

ARMSTRONG, C., AND J. DUFFY (2010): “Status and Coordination in World of Warcraft,” in preparation.

- COOPER, R. W., D. V. DEJONG, R. FORSYTHE, AND T. W. ROSS (1993): “Forward Induction in the Battle-of-the-Sexes Games,” *American Economic Review*, 83, 1303–1316.
- DEKEL, E., J. C. ELY, AND O. YILANKAYA (2007): “Evolution of Preferences,” *Review of Economic Studies*, 74, 685–704.
- HARSANYI, J. C. (1973): “Games with randomly disturbed payoffs: a new rationale for mixed-strategy equilibrium points,” *International Journal of Game Theory*, 2, 1–23.
- NASH, J. (1951): “Non-Cooperative Games,” *Annals of Mathematics*, 52, 286–295.