Incomplete Information in Repeated Coordination Games

January 24, 2017

Eric Rasmusen

Abstract

Asymmetric information can help achieve an efficient equilibrium in repeated coordination games. If there is a small probability that one player can play only one of a continuum of moves, that player can pretend to be of the constrained type and other players will coordinate with him. This hurts efficiency in the repeated battle of the sexes, however, by knocking out the pure-strategy equilibria.

Dan R. and Catherine M. Dalton Professor, Department of Business Economics and Public Policy, Kelley School of Business, Indiana University. Visitor (2007-2008), Nuffield College, Oxford University. Office: 011-44-1865 554-163 or (01865) 554-163. Nuffield College, Room C3, New Road, Oxford, England, OX1 1NF. Erasmuse@indiana.edu. http://www.rasmusen.org. This paper: Http://www.rasmusen.org/papers/goodtype-rasmusen.pdf. It is forthcoming in *Essays in Honor of Moriki Hosoe*, eds. Woohyung Lee, Tohru Naito, and Yasunori Ouchida, Springer.

I thank the Indiana Business Economics Brown Bag Lunch and Michael Baye for helpful comments.

1.Introduction

It is well known that coordination games have multiple equilibria, depending on player expectations, even if one equilibrium is pareto superior and players can communicate. This multiplicity is present even in the one- shot game, and just gets worse when the game is repeated. Few of the refinements of Nash equilibrium that have been suggested in the context of other kinds of games help with coordination games, and none has gained more than minimal acceptance.

The problem of multiple equilibria in coordination games has attracted attention from various authors. One way to try to predict which equilibrium is played out is to use the behavioral idea of "focal points" from Thomas Schelling (1960): that a human's attention is drawn to certain equilibria because they look "different". Thus, if a game's equilibria had payoffs of (1,1), (2,2), and (100,100), the focal point would be (100,100). This is a difficult notion to formalize, though: if the alternatives were (1,1), (99,99), and (100,100) would we predict that that the players would end up at (1,1) because it is the most distinctive?

Clearly, the idea of the focal point is important. Philosopher David Lewis (1965) divides the idea of the salience of a choice into two parts. The choice has "primary salience" to a player if he believes it is salient to he himself; it has "secondary salience" if he believes it has salience to other players. Judith Mehta, Chris Starmer & Robert Sugden (1994, p. 661) add "Schelling salience" to primary and secondary salience, as a choice that "seems obvious or natural to people who are looking for ways of solving coordination problems." In their article they report results of experiments trying to distinguish between primary salience– the answers subjects gave to questions when there was no reward for coordination– and secondary or Schelling salience– the answers when subjects were rewarded for successful coordination. They found that subjects indeed were picking with an eye towards what other subjects would pick; for example when asked to write down any day of the year, only 6% of the first set of subjects answered December 25, but 44% did when they were rewarded for successful cooperation.

A second approach tries to use derive the unique equilibrium from rationality. David Gauthier (1975, p. 201) defines the "Principle of Coordination" as "in a situation with one and only one outcome which is both optimal and a best equilibrium, if each person takes every other person to be rational and to share a common conception of the situation, it is rational for each person to perform that action which has the best equilibrium as one of the possible outcomes." Bacharach (1993) John Harsanyi & Reinhard Selten (1988) and Maarten Janssen (2000, 2001) have pursued this approach, trying to add axioms for rational behavior that require players to avoid dominated equilibria.

Repeating the game does not reduce the number of equilibria, but it does introduce a new angle: finding the optimal way to play a game starting without a convention as to the equilibria. What is the optimal strategy for the two players if they must first grope their ways towards coordination by guessing what the other player will do before they end up at the same action and use it thereafter? That is the project in Vincent Crawford & Hans Haller (1990), who find a learning procedure that converges in finite time.

A third approach is to look at evolution in games. Glenn Ellison (1993), Michihiro Kandori, George J. Mailath & Rafael Rob (1993), and Peyton Young (1993) take this approach. Start with a population of pairwise- interacting players with different strategies. They play coordination games, and increase or diminish in frequency depending on their payoffs. In such settings, "risk-dominant strategies" emerge as equilibria. In a symmetric two-player setting, this is the strategy a player would choose if he thought there was a 50% probability of the other player choosing each strategy. The risk-dominant strategy is not necessarily the one with the highest payoff; it balances that against the loss if discoordination does occur.

Risk-dominant equilibria also arise in the single-repetition "global games" of Stephen Morris & Hyun Song Shin (2003). They ask what happens if players have some small uncertainty over what game they are playing out. It turns out that iterated deletion of interim-dominated strategies can then make the risk-dominant equilibrium the unique equilibrium.

I will show below that adding incomplete information changes the repeated game drastically. David Kreps, Paul Milgrom, John Roberts & Robert Wilson (1982) show that adding a small amount of carefully chosen incomplete information to the model can result in cooperation in the finitely repeated prisoners' dilemma. Drew Fudenberg & Eric Maskin (1986) show more generally that adding incomplete information can generate any of a wide range of average payoffs in finite repeated games by getting around the backwards induction of the Chainstore Paradox. Their theorem does not apply to many coordination games, since it depends on a "dimensionality condition" that requires payoffs to vary enough between players to allow equilibria to be supported by punishment phases in which one player is able to punish another without hurting himself. Jean-Pierre Benoit & Vijay Krishna (1985), however, show that if a game has multiple equilibria, as a coordination game does, then a wide range of equilibria can be obtained if the game is repeated enough times by using the threat of punishment phase in an inferior equilibrium to enforce the desired behavior.

I will not be be able to reduce the number of equilibria in the one- shot game, but I will show that with a small amount of incomplete information and enough repetitions any perfect bayesian equilibrim of even a finitely repeated two-player game will achieve arbitrarily close to the optimal average payoff.

The results will not depend on careful specification of the incomplete information, and it is robust to out-of-equilibrium beliefs. There will be no assumption that "Players are either of type $x = 0$ or type $x = 100$ (with small probability), but never any other value of x." Nor will I specify anything like, "Out of equilibrium, the deviating player is believed to be of type $x \in [0, 34]$." Rather, the intuition is that in coordination games, no player has an incentive to hurt other players, so any attempt to "fool" other players by pretending to be of a particular type will be eagerly accepted by them. This intuition is partly present in the intuition behind the Gang of Four Theorem of Kreps, Milgrom, Roberts & Wilson (1982); here, it applies better and so the result is easier to achieve.

2. The Coordination Game with Complete Information

Consider a ranked coordination game with $n = 2$ players indexed by i who simultaneously choose actions x_1, x_2 from the interval [0, 100]. The per-period payoff to player i is $\pi(x_i, x_{-i})$ with:

$$(a)\forall x,\ \frac{\partial \pi(x,x)}{\partial x} > 0 \quad (b)\pi(0,0) > \pi(x_i, x_{-i})\ if\ x_i \neq x_{-i} \tag{1}$$

Assumption (a) says that a player's payoff rises if he chooses a higher action and the other player chooses the same action as he does. Assumption (b) says that if the players choose different actions their payoffs are lower than if they coordinated on (0,0).

We will normalize to $\pi(0,0) = 0$ and $\pi(100, 100) = 100$, which is to say the per-period payoff is 0 when both players choose $x = 0$ and 100 when they both pick $x = 100$. The assumptions then imply that coordination on $x > 0$ yields positive payoffs and discoordination yields negative payoffs.

If the game is unrepeated and $T = 1$ it has a continuum of pure strategy equilibria with x on the continuum from 0 to 100, as well as mixed strategy equilibria. All players prefer the equilibrium in which $x = 100$.

Which equilibrium will be played out depends on player expectations. A reasonable prediction is $x = 100$ because it is pareto superior to all other equilibria, a focal point. An equally special equilibrium, however, is $x = 0$. It is easy to imagine how the players could be caught in any equilibrium— if the game were preceded by a malicious outsider's cheap talk announcement that he expected them all to choose $x = 1$, for example, or if the players had a history of playing $x = 5$ for many periods.

Next, let the game be repeated a possibly infinite number T times, with the players observing each other's strategies after each round and no discounting. The equilibrium outcomes and strategies both become more numerous. Let us classify them as follows:

In a **time-dependent equilibrium**, some player's strategy in a round depends on which round number it is. If the strategies are the same in each round, the equilibrium is **time-independent**.

In a **history-dependent equilibrium**, some player's strategy in a round depends on the history of play up to that point. If the strategies do not depend on past play, the equilibrium is **history- independent**.

Table 1 provides examples.

		Time	
		Independent	Dependent
History	Independent	(a) Play 10 in each round.	(b) Play 20 in the first round and 25 in the second.
	Dependent	(c) Play 30 in each round unless someone deviates, in which case play 30 in the second round.	(d) For the first 50 rounds, player 1 picks 2 and the other players pick 14. For the last 10,000 rounds everyone picks 100 unless someone deviates. If someone deviates, all pick 0 for the remainder of the game.

Table 1: Four Types of Games

Note that the history-dependent equilibria include equilibria in which the players discoordinate in some periods, receiving flow payoffs of zero. Benoit & Krishna (1985) show that a wide array of outcomes might be observed in equilibrium, supported by punishment strategies similar to strategy (d) in table 1. The players choose any specified pattern of actions in the first S periods because in equilibrium they all play $x = 100$ in the last $(T - S)$ periods but if anybody deviates earlier they all play $x = 2$. The observed actions, for example, might be (10, 2), (7, 7), (8, 3), and then (100,100) for the last 200 periods. Thus, mere repetition of the game does not solve the problem of multiple equilibria, and in fact, even more outcomes become possible. The average payoff could even be negative, if the equilibrium has many periods of discoordination, so long as the average payoff is not below the discoordination payoff.

3. Incomplete Information: The Single-Action Player

Let us modify the game in the spirit of Kreps at al. (1982), by adding a small amount of incomplete information. Players are of two types. With some arbitrarily small probability $p > 0$, a player i must play $x_i = z_i$ in every round of the game, where z_i is chosen from [0, 100] using a atomless density $f(z_i)$ such that $f(100) > 0$. Such a player is "constrained"; otherwise, the player is "free". Note that the other players do not observe which players are constrained, their exact types, or even how many there are.

What is essential is that there be some possibility a player will choose $x = 100$ and stick with it, which is true of the specification above. All that is needed is a possibility, however small, and it can even have probability zero in the mathematical sense. That is the case in our specification, since any particular value of z has zero probability, despite having positive probability density. What that means is simply that we would predict any particular value of z (e.g. 97.345) with zero probability, even though we would predict a positive probability for any *interval* of types (e. g. [97, 98.5]).

How do we interpret the incomplete information? It might be that a constrained player is truly constrained, or that he misunderstands the rules of the game, or he is irrational and thinks all players will make the same choice as he does (psychology's "magical thinking"; see Brendan Daley & Philipp Sadowski [2016]). If we use a different specification, such as that there is a .0001% probability that a player is constrained to use $z = 100$, then we could interpret it as that the constrained player wishfully thinks that the equilibrium will be the pareto-optimal one (perhaps having read some of the references above) or thinks, for whatever reason, that if he starts with $x = 100$ the other players will join him.

In the modified game, some equilibria disappear, as Example 1 shows.

Example 1. Suppose $T = 20$ and the payoff from discoordination is -500. Is it an equilibrium for a free player to follow the strategy $x = 5$ in every period and for a constrained player of type z to play $x = z$? No.

Consider what happens if player 1 deviates to $x = 100$ in the first round. Is it a best response for player 2 to play $x = 5$ in the second round? That depends on player 2's beliefs,which are generated by Bayes's Rule:

$$Prob(z_1 = 100|x_1 = 100) = \frac{Prob(x_1=100|z_1=100)*Density(z_1=100)}{Prob(x_1=100|z_1=100)*Density(z_1=100)+Prob(x_1=100|z_1=free)*Prob(z_1=free)} \quad (2)$$

The priors tell us that $Prob(player\ 2\ is\ free) = 1 - p$ and $Density(z_1 = 100) = f(100)p$. In the proposed equilibrium, $Prob(x_1 = 100|z_1 = 100) = 1$ and $Prob(x_1 = 100|z_1 = free) = 0$. Thus, equation (2) becomes

$$Prob(z_1 = 100|x_1 = 100) = \frac{(1) * f(100)p}{(1) * f(100)p + (0) * (1 - p)} = 1. \quad (3)$$

After the first round, Player 2 therefore believes that Player 1's type is $z_1 = 100$, so he concludes that $x_1 = 100$ for all future rounds. Player 2's best response is not $x = 5$, but to imitate Player 1's action, deviating to $x_2 = 100$. If both players then stick with $x = 100$, their payoffs are $(-500 + 19(100), -500 + 19(100))$ compared to the $(20(5), 20(5))$ they would have gotten in the proposed equilibrium. Thus, Player 1's deviation has been profitable.

It is not true, however, that the only equilibrium in Example 1 is for a player to start with $x = 100$ and to choose in the second and succeeding periods whatever the other player chose in the first period. If $x = 99.9$, it is not worth bearing the initial cost of -500 to deviate. Rather, what we can say is that for large enough w a time- independent equilibrium strategy must have a player beginning with $x = w$ and then choosing in the second and succeeding periods whatever the other player chose in the first period. In such an equilibrium, the equilibrium payoff is $(20w, 20w)$. The optimal deviation is to $x = 100$, which generates a deviation payoff of $(-500 + 19(100), -500 + 19(100))$. There is no incentive to deviate from equilibrium if and only if $w \geq 92.5$.

Example 1 is the essence of this paper. If information is incomplete, then a player can break out of a bad equilibrium at some cost by pretending to be of an unusual type. If the game is repeated long enough, it is worthwhile to bear that cost. Thus, if T is large enough, the game has a much smaller interval of equilibria and the average payoff becomes arbitrarily close to 100.

Proposition 1: *For any ϵ, there exists T large enough that in all pure-strategy equilibria the average payoff approaches within ϵ of the optimum:*

$$\forall \epsilon > 0, \exists T : \frac{\sum_{t=1}^{T} \pi_{it}}{T} > 100 - \epsilon. \tag{4}$$

Proof. The probability that a player is constrained is an arbitrarily small p, so the effect that the presence of truly constrained players have on the average equilibrium payoffs will be less than ϵ.

Let the equilibrium with the lowest average payoff call for the players to first choose (a, b) with a or b or both not equal to 100 in round t_1. Without loss of generality, suppose that player 1 chooses $a \neq 100$.

The minimum bound on the payoff is set by player 1 having the deviation option to choose $x = 100$ in that period and convince player 2 that player 1 is constrained of type $z = 100$. Both players would choose $x = 100$ for every succeeding round. This would generate a payoff of $\pi(100, b) + 100(T - 1)$, where $\pi(100, b)$ is the discoordination payoff that arises from that particular deviation, since there would be one period of discoordination and all other periods will have per-period payoffs of 100. This strategy will have an average payoff of

$$\frac{\pi(100, b)}{T} + \frac{100(T-1)}{T} = 100 + \frac{\pi(100, b)}{T} - \frac{100}{T}. \tag{5}$$

If T is large enough, the last two terms, which are both negative, shrink to less than whatever small amount ϵ we might choose. Q.E.D.

The equilibria will be in actions with an average payoff in the interval $[100 - \epsilon, 100]$ for some ϵ that depends on T. This set of equilibria does not depend heavily on the out-of-equilibrium payoffs–just for one period of discoordination loss– and therefore it is not necessarily the same as the set of risk-dominant equilibria. It could be, for example, that for x in [0,50] the discoordination payoff if the other player chooses a different x is -1, but for x in (50, 100] it is $-5,000$, in which case the risk-dominant equilibrium would be (50,50), not (100,100).

4. Three or More Players in the Incomplete Information Game

Now let us allow for more than two players. Consider a ranked coordination game with $n \geq 2$ players indexed by i who simultaneously choose actions $x_1, ..., x_n$ from the interval [0,100]. If $m(x_i)$ players choose the same action x_i, the per-period payoff to player i is $\pi_i(x_i, x_{-i}, m(x_i))$, with:

$$
\begin{aligned}
&(a')\frac{\partial \pi_i(x_i,x_{-i},m(x_i))}{\partial x_i} \geq 0 \\
&(b')\frac{\Delta \pi_i(x_i,x_{-i},m(x_i))}{\Delta m(x_i)} > 0, \\
&(c)\frac{\partial^2 \pi_i(x_i,x_{-i},m(x_i))}{\partial x_i \partial x_j} = 0 \\
&(d)\pi_i(0, x_{-i}, n) > \pi_i(100, x_{-i}, n-1), \\
&(e)\pi_i(w, x_{-i}, l) > \pi_i(w', x_{-i}, l-1) \forall l, w, w' \neq x
\end{aligned}
\tag{6}
$$

Assumption (a') says that the payoff to player i rises or stays the same as the magnitude of the group action x_i rises– from 88, say, to 89. Assumption (b') says that the payoff to choosing action x_i rises from being in a bigger group.

Assumption (c) says that the payoff to player i from choices made by players who choose discoordinating actions does not depend on which actions they choose.

Assumption (d) says that group size matters more than action size: the payoff to i from choosing $x_i = 0$ in a group of n is bigger than from choosing $x_i = 100$ in a group of size $n-1$. Assumption (e) is a more general version of (d), saying that a larger group always gets a bigger payoff, no matter what the size of the action.

My colleagues Michael Rauh and Michael Baye suggested the following as a payoff function that satisfies assumptions (a') through (e)

$$\pi_i(x_i, x_{-i}, m(x_i)) = [m(x_i)(1 + \tfrac{x_i}{1000}) - n](100/1.1n) \tag{7}$$

or, without our normalization of $\pi_i(0, x_{-i}, n) = 0$ and $\pi_i(100, x_{-i}, n) = 100$,

$$\pi_i(x_i, x_{-i}, m) = m(x_i)(1 + \frac{x_i}{1000}) \tag{8}$$

The complete information game has the usual continuum of equilibria, just as when there are just two players. How about the incomplete information game? Consider Examples 2 and 3.

Example 2. Let there be incomplete information of the following form: with some arbitrarily small probability $p > 0$, player i is "constrained" and must play $x_i = z_i$ in every round of the game, where z_i is chosen from [0, 100] using a atomless density $f(z_i)$ such that $f(100) > 0$. Let there be three players, and consider whether it is an equilibrium outcome to play (5, 5, 5) each of T periods. Suppose player 1 deviates to $x_1 = 10$ in the first period. A unilateral switch by one of the other two players from 5 to 10 would be profitable in the long run in the incomplete information game.

Example 3. Let there be incomplete information of the following form: with some arbitrarily small probability $p > 0$, player i is "constrained" and must play $x_i = z_i$ in every round of the game, where z_i is chosen from [0, 100] using a atomless density $f(z_i)$ such that $f(100) > 0$. Let there be four players, and consider whether it is an equilibrium outcome to play (5, 5, 5, 5) each of T periods. Suppose player 1 deviates to $x_1 = 100$ in the first period. Even if player 1 were truly of type $z = 100$, a unilateral switch by one of the other three players from 5 to 100 would be unprofitable if $\pi_i(x = 5, m = 3) > \pi_i(x = 100, m = 2)$. Thus, (5, 5, 5, 5) would be an equilibrium outcome. Incomplete information does not reduce the number of equilibria.

The situation changes if we change the form of the incomplete information. What is needed now is a more-than-infinitesimal probability of a given constraint value of z. Recall our alternative specification of incomplete information in which player i has probability .01 of being constrained to play $x_i = 100$. In that specification, the probability is not zero; it is strictly positive.That makes a huge difference, as we see in Example 4.

Example 4. Let there be incomplete information of the following form: with some small probability $p > 0$, player i is "constrained" and must play $x_i = 100$ in every round of the game. Let there be four players, and consider whether it is an equilibrium outcome to play (5, 5, 5, 5) each of T periods. Assume that if the players do not all play the same action, their period payoff is zero. Suppose player 1 deviates to $x_1 = 100$ in the first period. With probability $.99^3$, he is the only player to play 10, and the other players do not imitate him in future periods for the same reason as in Example 3. But with probability $1 - .99^3$, at least one of the other players is a true constrained player who also plays 10. If that happens, then in the second and succeeding periods, the remaining two players will play 10. Thus, the expected payoff from deviating (and returning to playing 5 if no other player player 10 in the first period) will be greater than $.99^3 * (0 + (T-1)(5)) + (1 - .99^3)(0 + (T-1)(100)$ (I say "greater than" because it is slightly higher because of the possibility that not just one, but two or even all three of the other players are constrained to play 100). If T is great enough, the deviation payoff is greater than the proposed equilibrium payoff of $5T$. It is worth the high probability of one period with a payoff of 0 in order to have a chance at $(T - 1)$ periods with a payoff of 100. If the deviation is profitable, however, then all four of the players will choose 100 in the first and every period, constrained or not.

Thus, with four players, or more, incomplete information can still justify a unique efficient equilibrium. The story is a little different, though, because if expectations begin with some action less than 100, the player who deviates puts high probability on his deviation being unprofitable— it is just that if it does work out successfully he gets a very large payoff increase. For this to work, T must be much larger than when there are only two or three players.

5. Mixed Coordination-Conflict Games: The Battle of the Sexes

Incomplete information can actually hurt in mixed coordination- conflict games, by destroying the possibility of pure-strategy equilibria. Consider the Battle of the Sexes in Table 2. It has two pure strategy equilibria, $(prizefight, prizefight)$ and $(ballet, ballet)$, and a mixed strategy equilibrium. in which the man plays $prizefight$ with probability $m = A/(A + B)$ and the woman with probability $w = B/(A + B)$.

The total payoff in the two pure-strategy equilibria are $(A,)$ and (B, A). The man's one-shot expected payoff is then $AB/(A + B)$, which is less than B since $B < A + B$. The man's expected payoff (and analogously the woman's) in the mixed- strategy equilibrium is lower even than in the pure-strategy equilibrium he likes least.

TABLE 2: THE BATTLE OF THE SEXES

		Woman	
		prizefight	*Ballet*
Man	*prizefight*	**A,B**	0,0
	Ballet	0,0	**B,A**

Payoffs to: (Man, Woman). $A > B$.

The T-repeated game has many subgame perfect equilibria, but let us focus on the three time-independent and history-independent equilibria that repeat the single-period equilibria just described.

Now let us add incomplete information, in the form of constrained players. With probability p_1, the man is constrained to play $Prizefight$ and with independent probability p_2 the woman is constrained to play $Ballet$. This change eliminates $(Ballet, Ballet)$ and $(Prizefight, Prizefight)$ as equilibria. Suppose the woman thought $(Ballet, Ballet)$ was the equilibrium. The man would begin the game by playing $Prizefight$. The woman would conclude that the man was constrained, and would play $Prizefight$ in all future rounds, so the man would have succeeded in increasing his payoff (if there are enough rounds) to one round of (Ballet, Prizefight) and $T - 1$ rounds of $(Prizefight, Prizefight)$. The equilibrium $(prizefight, prizefight)$ would similarly fail.

The mixed strategy equilibrium survives. If the man deviates to playing $prizefight$ as a pure strategy, the woman will interpret this as a realization of the equilibrium strategy. This is ironic, however, because the man's ability to knock out the pure-strategy equilibrium of $(Ballet, Ballet)$ ends up hurting him: his payoff is higher in that equilibrium than in the mixed-strategy equilibrium that survives.

Closing Remarks

Thus, we see that in repeated Ranked Coordination, the efficient equilibria are robust to incomplete information but the inefficient equilibria are not, whereas in the Battle of the Sexes the opposite is true. This model has used a particular specification of incomplete information, to be sure, but if we added other incomplete information without removing the possibility of constraint in this model, the results would often stay the same. The key to the result is that if there is some chance that a deviation beneficial to the deviator will be interpreted as predicting that he will choose the same action in the future, the player will have incentive to deviate.

References

Bacharach, Michael (1993) "Variable Universe Games," *Frontiers of Game Theory*, 1993.

Benoit, Jean-Pierre & Vijay Krishna (1985) "Finitely Repeated Games," *Econometrica,* 53: 905-922.

Binmore, Ken & Larry Samuelson (2006) "The Evolution of Focal Points," *Games and Economic Behavior*, 55: 21-42.

Daley, Brendan & Philipp Sadowski (2016) "Magical Thinking: A Representation Result," forthcoming, *Theoretical Economics.*

Carlsson, Hans & Eric van Damme (1994) "Global Games and Equilibrium Selection Hans Carlsson," *Econometrica*, 61: 989-1018.

Crawford, Vincent P. & Hans Haller (1990) "Learning How to Cooperate: Optimal Play in Repeated Coordination Games," *Econometrica,* 58: 571-595.

Glenn Ellison (1993) "Learning, Local Interaction, and Coordination," *Econometrica,* 61: 1047-1071.

Fudenberg, Drew & Eric Maskin (1986) "The Folk Theorem in Repeated Games with Discounting or with Incomplete Information," *Econometrica*, 54: 533-554.

Gautier, David (1975) *Coordination*, Cambridge University Press.

Goyal, Sanjeev & Maarten C. W. Janssen (1996) "Can We Rationally Learn to Coordinate?" *Theory and Decision,* 40: 29-49.

Goyal, Sanjeev & Maarten C. W. Janssen (1997) "Non-Exclusive Conventions and Social Coordination," *Journal of Economic Theory*, 77: 34–57 (1997).

Janssen, Maarten C. W. (2001) "On the Principle of Coordination," *Economics and Philosophy* 17: 221-234.

Janssen, Maarten C. W. (2001) "Rationalizing Focal Points," *Theory and Decision,* 50: 119-148.

Kandori, Michihiro, George J. Mailath & Rafael Rob (1993) "Learning, Mutation, and Long Run Equilibria in Games," *Econometrica*, 61: 29-56.

Kreps, David, Paul Milgrom, John Roberts, & Robert Wilson (1982) "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma," *Journal of Economic Theory*, 27: 245-252.

Lewis, David (1969) *Convention,* Harvard University Press, Cambridge, Mass.

Mehta, Judith, Chris Starmer & Robert Sugden (1994) "Focal Points in Pure Coordination Games: An Experimental Investigation," *Theory and Decision*, 36: 163-185.

Morris, Stephen & Hyun Song Shin (2003) "Global Games: Theory and Applications," in *Advances in Economics and Econometrics (Proceedings of the Eighth World Congress of the Econometric Society)*, edited by M. Dewatripont, L. Hansen and S. Turnovsky. Cambridge, England: Cambridge University Press (2003).

Schelling, Thomas (1960) *The Strategy of Conflict*.

Sugden, Robert (1995) "A Theory of Focal Points," *The Economic Journal,* 105: 533-550.

Young, H. Peyton (1993) "The Evolution of Conventions," *Econometrica*, 61: 57-84.