Coordination and Higher Order Uncertainty

Jonathan Levin

April 2006

In these notes, we discuss work on coordination in situations of uncertainty and investigate the importance of "higher order beliefs" — that is, players' beliefs about their opponents beliefs about their beliefs about.

We start with a particular class of "global" coordination games of incomplete information where players' beliefs are highly, but not perfectly, correlated. There games are interesting for several reasons. First, they capture in simple form the idea that in strategic settings where actions are conditioned on beliefs, in particular settings where coordination is important, players need to be concerned with what their opponents believe, what their opponents believe about their beliefs, and so on. Second, global games can allow us to refine equilibria in coordination games in a very strong way. In some models, we can use global games analysis to show that even if common knowledge of payoffs gives rise to multiple equilibria, there will be a unique equilibrium if the players' information is perturbed in even a "small" way. A recent applied literature has arisen using these techniques particularly in finance and macroeconomic (Morris and Shin, 2002, is a nice survey).

We then use related ideas to look at the difference between public and private information in situations where higher order beliefs matter greatly. Finally, we discuss the broader relevance of higher order beliefs in "typical" games. We discuss the "types" approach of Harsanyi, the modeling of higher-order uncertainty, and some further applications.

1 Global Games

We work with an example drawn from Morris and Shin (2002) and based on Carlsson and van Damme (1993). There are two players $i = 1, 2$ who choose one of two actions, "Invest" or "Not Invest". The payoffs are as follows:

	Invest	Not Invest
Invest	θ, θ	$\theta - 1, 0$
Not Invest	$0, \theta - 1$	$0, 0$

Not Invest is a safe action that yields zero, while Invest yields $\theta - 1$ if the opponent doesn't invest and θ if she does. If θ is known to the players, there are three possibilities:

- $\theta > 1$, Invest is a dominant strategy
- $\theta \in [0, 1]$, (Invest, Invest) and (NI,NI) are both NE.
- $\theta < 0$, Not invest is a dominant strategy

Suppose we introduce incomplete information by assuming that each player does not observe θ, but rather a private signal $x_i = \theta + \varepsilon_i$, where $\varepsilon_i \sim N(0, \sigma^2)$, each ε_i is independent. Suppose also that θ is drawn from a uniform distribution over the entire real line (i.e. players have an uninformative or improper prior on θ).

Proposition 1 *In the incomplete information game, there is a unique equilibrium in which both players invest if and only if* $x_i > x^* = \frac{1}{2}$.

Proof. Let's first verify that the stated equilibrium is in fact an equilibrium. First, observe that if player j uses the stated strategy, then i's payoff from investing conditional on having signal x_i is

$$\mathbb{E}[\theta \mid x_i] - \Pr\left[x_j \leq \frac{1}{2} \mid x_i\right] = x_i - \Phi\left(\frac{\frac{1}{2} - x_i}{\sqrt{2}\sigma}\right),$$

where we use the fact that $x_j | x_i \sim N(x_i, 2\sigma^2)$. This payoff increasing in x_i, and equal to zero when $x_i = 1/2$. So i's best response is to invest if and only if $x_i \geq 1/2$, verifying the equilibrium.

To prove uniqueness, we'll show that the stated profile is the only one that survives iterated deletion of dominated strategies. As a preliminary, consider i's payoff from investing conditional on having signal x_i and conditional on j playing the strategy "invest if and only if $x_j \geq k$" for some "switching point" k. The payoff is

$$\mathbb{E}[\theta \mid x_i] - \Pr[x_j \leq k \mid x_i] = x_i - \Phi\left(\frac{k - x_i}{\sqrt{2}\sigma}\right),$$

which is again increasing in x_i, and also decreasing in k. Let $b(k)$ be the unique value of x_i at which the payoff to investing is zero. Observe that $b(0) > 0$, $b(1) < 1$, that $b(\cdot)$ is strictly increasing in k, and that there is a unique value k^* that solves $b(k) = k$, namely $k^* = 1/2$.

We now show that if a strategy s_i survives n rounds of iterated deletion of strictly dominated strategies, then:

$$s_i(x) = \begin{cases} \text{Invest} & \text{if } x_i > b^{n-1}(1) \\ \text{Not Invest} & \text{if } x_i \leq b^{n-1}(0) \end{cases}.$$

Note that the value of $s_i(x)$ for values of x between $b^{n-1}(0)$ and $b^{n-1}(1)$ is not pinned down here.

This claim follows from an induction argument. Suppose $n = 1$. The worst case for i investing is that j never invests (i.e. uses a switching strategy with $k = \infty$). If j never invests, then i should invest if and only if $\mathbb{E}[\theta|x_i] > 1$, or in other words if $x_i \geq 1$ (recall that $\mathbb{E}[\theta|x_i] = x_i$). This means that Not Invest is dominated by Invest if and only if $x_i > 1$. Conversely, Invest is dominated by Not Invest if and only if $x_i < 0$. This verifies the claim for $n = 1$.

Now suppose the claim is true for n. Now the worst case for i investing is that j invests only if $x_j > b^{n-1}(1)$, i.e. plays a switching strategy with cut-off $k = b^{n-1}(1)$. In this case, i should invest if and only if $x_i > b(b^{n-1}(1)) = b^n(1)$, meaning that Not Invest is in fact dominated if $x_i > b^n(1)$. Conversely, Invest is dominated by Not Invest at this round if and only if $x_i < b(b^{n-1}(0)) = b^n(0)$. So we have proved the claim.

To complete the result, we observe that:

$$\lim_{n\to\infty} b^n(0) = \lim_{n\to\infty} b^n(1) = \frac{1}{2}.$$

Therefore iterated deletion of dominated strategies yields a unique profile in which both players invest if and only if their respective signal exceeds one-half. *Q.E.D.*

1.1 Generalizing the Model

The logic of this example can be generalized without much trouble to other two-action two-player supermodular games. In particular, suppose we have two players, each of who chooses an action $a \in \{0, 1\}$. For simplicity, they symmetric payoffs $u : \{0, 1\} \times \{0, 1\} \to \mathbb{R}$, where $u(a_i, a_j, x_i)$ is i's payoff from choosing a_i, given that j chooses a_j, and that i's private signal is x_i.

Signals are generated in the following way. First, nature selects a state $\theta \in \mathbb{R}$ drawn from the (improper) uniform distribution on $\mathbb{R}$. Player i then observes a signal $x_i = \theta + \sigma\varepsilon_i$, where ε_i has a continuous density $f(\cdot)$ with support $\mathbb{R}$. Call this game $G(\sigma)$.

Define the incremental returns to choosing $a = 1$ as:

$$\Delta(a_j, x) = u(1, a_j, x) - u(0, a_j, x).$$

We impose the following assumptions.

1. **Action Monotonicity**. Δ is increasing in a_j.

2. **State Monotonicity**. Δ is strictly increasing in x.

3. **Continuity.** Δ is continuous in x.

4. **Limit Dominance**. There exist $\underline{\theta}, \overline{\theta} \in \mathbb{R}$ such that $\Delta(a_j, x) < 0$ whenever $x < \underline{\theta}$, and $\Delta(a_j, x) > 0$ whenever $x \geq \overline{\theta}$.

Proposition 2 *Under (A1)–(A4), the essentially unique strategy profile that survives iterated deletion of strictly dominated strategies in $G(\sigma)$ satisfies $s(x) = 0$ for all $x < \theta^*$ and $s(x) = 1$ for all $x > \theta^*$, where θ^* is the unique solution to:*

$$(\Delta(1, \theta^*) + \Delta(0, \theta^*))/2 = 0.$$

Proof. Nearly identical to the one above. *Q.E.D.*

In the unique equilibrium, each player uses a cut-off strategy (we say essentially unique because the strategy is indeterminate at $x = \theta^*$. Moreover, the cut-off θ^* is such that if i receives the signal θ^*, and believes that j is equally likely to play 0 or 1, then i will be just indifferent between playing 0 and 1.

Note one difference between this game and the first one is that this has "private values" while the other game had "common values". This distinction is not so important, however. Suppose payoffs in the more general case were $u(a, \theta)$, rather than $u(a, x)$. We could simply define $\Delta(a_j, x)$ as the *expected* return to playing 1 rather than 0 after observing a signal x, given that one's opponent was playing a_j. Assuming j was using a cut-off strategy, Δ would satisfy all the same properties, so everything will still go through.

1.2 Discussion

We now discuss several features and extensions of the model.

1. (Inefficiency of the Unique Equilibrium). If we consider a sequence of incomplete information games with $\sigma \to 0$, then, since in equilibrium each player invests if and only if $x = \theta + \varepsilon \geq \frac{1}{2}$, in the limit players coordinate on (Invest,Invest) whenever $\theta > \frac{1}{2}$, and on (Not Invest,Not Invest) whenever $\theta \leq \frac{1}{2}$. Coordination on $(Invest, Invest)$, however, is *efficient* whenever $\theta > 0$. So the fact that the players act in a decentralized fashion means that they generally won't coordinate efficiently.

2. (Risk Dominance) In 2×2 symmetric games, an action is *risk-dominant* if it is a best-response given that one's opponent is mixing uniformly. In the underlying symmetric information game, investing is risk-dominant if $\theta \geq \frac{1}{2}$ and not investing is risk-dominant if $\theta \leq \frac{1}{2}$. Hence as $\sigma \to 0$, and we converge to the complete information game, the players play the risk dominant action.

3. (Equilibrium Refinements) More generally, note that if $0 < \theta < \frac{1}{2}$, (Invest, Invest) is an equilibrium of the complete information game, but not of the closely related incomplete information game (with σ small but positive). This can be related to the general problem of selecting more or less plausible equilibria in a given game — or *refining* the set of equilibria. A common approach in this regard is to consider a family of nearby games and ask if these games have equilibria that are "close" to a given equilibrium of the original game. Kajii and Morris (1997) say that an equilibrium of a given game is *robust* if it is an equilibrium of all nearby games of incomplete information. Some games have no robust equilibria, but Kajii and Morris show that some interesting classes of games do have robust equilibria.

4. (Generalizations) It is quite easy to duplicate the above analysis to more general 2×2 games with strategic complementarities, provided some technical conditions are satisfied (see Morris and Shin, 2002). Frankel, Morris and Pauzner (2002) extend the above result to asymmetric n-player many action games with strategic complementarities. They provide conditions under which, if there is only a small amount of noise, equilibrium will be unique. The selected equilibrium, however, may depend on the fine structure of the noise. Interestingly, however, if a game has a unique robust equilibrium, this equilibrium will be

selected regardless of the noise structure. Frankel and Pauzner (2000) and Levin (2000) use global game arguments to identify unique equilibria in dynamic games with strategic complementarities. One such problem is on your homework.

5. (Higher Order Beliefs) One way to understand why the incomplete information game has a unique equilibrium when nature selects $\theta \in (0, 1)$ despite the complete information game having multiple strict Nash equilibria, is that there is a failure of *common knowledge* in the incomplete information game. We return to this point below.

6. (Applications) There has been a lot of interest in the application of global games to macroeconomics, in particular currency crises — the view being that currency pegs tend to fall when there is an attack by many investors, a situation that naturally gives rise to a coordination game. A twist in such settings is that the ability to observe prices may restore common knowledge (one analysis of this is by Angeletos and Werning, 2005).

2 Public and Private Information

In the coordination game above, players care about fundamentals (the value of θ) and also about the actions of their fellow player (the value of a_{-i}). Each player's signal is informative about both variables of interest. Of course, in the context of the above example, information is private. An interesting question that arises in a coordination setting concerns the role of *public* information. Intuitively, public information about fundamentals should be valuable in a coordination setting because (barring unfortunate coordination failure) it will allow for coordination on the more appropriate action. Morris and Shin (2003), however, show that this intuition fails in settings where players have private information as well as public information. The basic idea, as we will see shortly, is that players are "too sensitive" to the public information because even if it is not that informative about fundamentals, it tends to be quite informative about other players' beliefs and hence about their actions.

2.1 The "Beauty Contest" Model

Morris and Shin's model is based on a famous "beauty contest" parable in Keynes' *General Theory.*[1] There are a continuum of agents, indexed by $i \in [0,1]$. Agent i chooses a_i, and receives a payoff:

$$u_i(a_i, a_{-i}, \theta) = -(1-r)(a_i - \theta)^2 - r(L_i - \overline{L}),$$

where $0 < r < 1$ and:

$$L_i = \int_0^1 (a_j - a_i)^2 dj \qquad \text{and} \qquad \overline{L} = \int_0^1 L_j dj \ .$$

Agent i wants to minimize the distance between his action and the true state θ and also minimize the distance between his action and the actions of the other agents. The parameter r weights these two parts of the objective function.

The beauty contest part of the game has a zero-sum flavor. If we define social welfare as the average of individual utilities:

$$W(a, \theta) = \int_0^1 u_i(a, \theta) di = -(1-r) \int_0^1 (a_i - \theta)^2 di.$$

From a social point of view, what matters is how close the individual actions are to θ, not to each other.

Each agent will maximize his utility by choosing a_i to minimize his expected loss:

$$a_i = (1-r)\mathbb{E}_i[\theta] + r\mathbb{E}_i[\overline{a}],$$

where $\overline{a} = \int_0^1 a_j dj$ is the population average utility, and $\mathbb{E}_i$ is the expectation operator conditioning on i's information.

2.2 Public Information Benchmark

As a benchmark, consider the case where the agent's only have public information. Suppose everyone shares an improper uniform prior on θ and then observes a public signal:

$$y = \theta + \eta,$$

[1] Keynes drew an analogy between the stock market and a certain beauty contest in a London newspaper. The paper printed pictures of young women. Keynes said that making money in the stock market was like trying to pick the girl who the most people would vote for as most beatiful — what mattered was not the true beauty of the girls, but whether or not people would vote for them.

where $\eta \sim N(0, \sigma_\eta^2)$. Then $\mathbb{E}_i[\theta|y] = y$ by Bayesian updating, and the unique equilibrium has each agent choose:

$$a_i(y) = y.$$

In this equilibrium, the expected welfare is:

$$\mathbb{E}(W|\theta) = -(1-r)\mathbb{E}[(y-\theta)^2|\theta] = -(1-r)\sigma_\eta^2.$$

So clearly better public information unambiguously improves social welfare.

2.3 Public and Private Information

In contrast, now suppose that in addition to y, each agent observes a *private* signal:

$$x_i = \theta + \varepsilon_i,$$

where $\varepsilon_i \eta \sim N(0, \sigma_\varepsilon^2)$, and moreover, ε_i and ε_j are independent for all $i \neq j$. In this set-up, agent i's information is the pair (x_i, y). She needs to use this information to forecast both the true state θ and the average action in the population.

By the wonderful properties of Bayes updating with normal random variables, we have:

$$\mathbb{E}_i[\theta|x_i, y] = \frac{h_\eta y + h_\varepsilon x_i}{h_\eta + h_\varepsilon},$$

where $h_\varepsilon = 1/\sigma_\varepsilon^2$ and $h_\eta = 1/\sigma_\eta^2$ are the precisions of ε and η.

We look for a linear equilibrium of the form:

$$a_i(x_i, y) = \kappa x_i + (1-\kappa)y.$$

If the equilibrium has this linear form, then:

$$\begin{aligned}
\mathbb{E}_i[\overline{a}|x_i, y] &= \kappa\left(\frac{h_\eta y + h_\varepsilon x_i}{h_\eta + h_\varepsilon}\right) + (1-\kappa)y \\
&= \left(\frac{\kappa h_\varepsilon}{h_\eta + h_\varepsilon}\right) x_i + \left(1 - \frac{\kappa h_\varepsilon}{h_\eta + h_\varepsilon}\right) y.
\end{aligned}$$

Agent i's optimal action is:

$$\begin{aligned}
a_i(x_i, y) &= (1-r)\mathbb{E}_i[\theta] + r\mathbb{E}_i[\overline{a}] \\
&= \left(\frac{h_\varepsilon(r\kappa + 1 - r)}{h_\eta + h_\varepsilon}\right) x_i + \left(1 - \frac{h_\varepsilon(r\kappa + 1 - r)}{h_\eta + h_\varepsilon}\right) y,
\end{aligned}$$

meaning we have a linear equilibrium with:

$$\kappa = \frac{h_\varepsilon(1-r)}{h_\eta + h_\varepsilon(1-r)}.$$

In this equilibrium, agent i's action is:

$$a_i(x_i, y) = \frac{h_\eta y + h_\varepsilon(1-r)x_i}{h_\eta + h_\varepsilon(1-r)}.$$

Morris and Shin (2003) verify that this is the unique equilibrium in their model. You can consult their paper for what they describe as a "brute force" proof.

2.4 Discussion and Welfare Properties

The key point to notice about the equilibrium behavior in the beauty contest model is that agents actions *over-weight* public information relative to its informativeness about economic fundamentals. In particular, both $\mathbb{E}_i[\theta|x_i, y]$ and $a_i(x_i, y)$ are linear combinations of x_i and y. In forming her expectation of θ, agent i puts a weight $h_\eta/(h_\eta + h_\varepsilon)$ on y. But in choosing her action, agent i puts a weight $h_\eta/(h_\eta + (1-r)h_\varepsilon)$ on y. Why the larger weight? Because even if x_i and y were to be equally informative about θ ($h_\eta = h_\varepsilon$), the public signal y would be *more informative* about other player's beliefs, and hence about their actions. An early version of Morris and Shin's paper referred to this as the "publicity multiplier".

Because agent's are forecasting other agents' actions — and hence other agents' beliefs and beliefs about beliefs and so on — public information is given disproportionate weight relative to its true informativeness about fundamentals. This can give rise to surprising welfare effects.

In particular, suppose we re-write the equilibrium strategy of each agent i as:

$$a_i = \theta + \frac{h_\eta \eta + h_\varepsilon(1-r)\varepsilon_i}{h_\eta + h_\varepsilon(1-r)}.$$

Then expected welfare is given by:

$$\mathbb{E}[W|\theta] = -(1-r)\frac{h_\eta + h_\varepsilon(1-r)^2}{(h_\eta + h_\varepsilon(1-r))^2}.$$

An increase in h_ε, the informativeness of the private signals, has an unambiguously positive effect on social welfare. On the other hand, an increase in h_η, the informativeness of the public signal, has an effect:

$$\frac{\partial \mathbb{E}[W|\theta]}{\partial h_\eta} \overset{sign}{=} h_\eta - (2r-1)(1-r)h_\varepsilon.$$

Better public information is always beneficial is $r < 1/2$, so that the "beauty contest" incentive is relatively small. If the beauty contest component of payoffs is large, however, so that $r > 1/2$ and there is a significant element of coordination involved in the equilibrium, better public information improves welfare only if the public information is reasonably good relative to the quality of private information — if h_η is small relative to h_ε, an increase in h_η will cause the agents' too substitute toward y in choosing their actions and increase the variance in the population action around θ.

3 Modeling Beliefs and Higher Order Beliefs

In this section, we take a short detour to discuss the modeling of beliefs and higher order beliefs. A running theme in the coordination models we have looked at is the importance not only of the players' beliefs about underlying fundamentals (i.e. about the true payoffs), but also about the beliefs of others. As emphasized in the last model, information about the beliefs of others (so as to ascertain their actions) can easily be more important than beliefs about fundamentals.[2]

We have modeled beliefs, and beliefs about beliefs, and so on, as arising from a model where players have a common prior belief about some objective payoff uncertainty, and update based on a private signal about payoffs. A player's signal conveys information about fundamentals and about the signals of others. Indeed, the one signal determines a player's "first-order" beliefs about fundamentals, his second-order beliefs about the first-order beliefs of others, his third-order beliefs about the second-order beliefs others have about his first-order beliefs, and so on.

[2] Note that higher-order beliefs can also be relevant in *complete information* games, in the sense that player i has to form a belief about player j's action, j has to form a belief about i's belief about j's action, and so on. Nash equilibrium "cuts through" this infinite regress of beliefs by assuming the first order beliefs are correct and player are rational with respect to these correct beliefs. Rationalizability, on the other hand, makes use of higher order beliefs to eliminate strategies that only iteratively dominated. With incomplete information, there is a sense in which beliefs about beliefs must be tackled head on — because to the extent that j has private information about i's payoff, i cares *directly* about j's belief, as well as about j's action.

This formulation, while powerful, is also limiting in the sense that a player's beliefs about fundamentals uniquely tie down his higher-order beliefs. A natural question is whether it is possible to develop models with a greater diversity of higher-order uncertainty. It certainly seems possible to posit a model where we specify first the beliefs of the players about fundamentals, then their beliefs about each other's beliefs, and so on. Harsanyi (1968) argued that rather than write down an entire hierarchy of beliefs, it would be possible to capture even higher order uncertainty in a model similar to the ones with which you are familiar. That is, Harsanyi argued that one could consider a model where each player initially drew a random "type" that described *all* of his beliefs.

Mertens and Zamir (1985) and Brandenburger and Dekel (1993) made this idea precise. In doing so, they show that it is possible to define a "universal" type space that would allow for all possible hierarchies of belief. These papers are quite technical, particularly Mertens and Zamir's, but it is worth sketching the ideas. I'll follow Brandenburger and Dekel, glossing over virtually all of the technical issues.

The starting point is a space of possible fundamentals S (e.g. payoffs). Each player will have a first-order belief about fundamentals, that is some $t_1 \in T_1 = \Delta(S)$. Each player will then have a second-order belief over fundamentals and the other players first order belief, that is some $t_2 \in T_2 = \Delta(S \times \Delta(S))$. Formally, define the spaces:

$$\begin{aligned} X_0 &= S \\ X_1 &= X_0 \times \Delta(X_0) \\ &\vdots \\ X_n &= X_{n-1} \times \Delta(X_{n-1}) \end{aligned}$$

The space of each player's nth order beliefs is then $T_n = \Delta(X_{n-1})$. A *type* for player i is a hierarchy of beliefs $t^i = (t_1^i, t_2^i, ...) \in \times_{n=1}^{\infty} T_n$. Let $T = \times_{n=1}^{\infty} T_n$ denote the space of all possible types for a given player.

Under this formulation, i knows his own type but not the type of his opponent j. So perhaps we will need to specify a belief for i about j's type, a belief for j about i's belief and so forth. Brandenburger and Dekel show that one additional assumption, however, pins down i's belief about j's type.

Definition 1 *A type $t = (t_1, t_2, ...) \in T$ is **coherent** if for every $n \geq 2$, $marg_{X_{n-2}} t_n = t_{n-1}$, where $marg_{X_{n-2}}$ denotes the marginal probability distribution on the space X_{n-2}.*

This simply says that i's beliefs do not contradict one another. This leads to the following result, for which you will need to know that a *homeomorphism* is a 1-1 map that is continuous and has a continuous inverse. For this result, let T' denote the set of all coherent types.

Proposition 3 *There is a homeomorphism* $f : T' \to \Delta(S \times T)$*.*

This result says that i's hierarchy of beliefs (his "type") also determines i's belief about j's type. Brandenburger and Dekel go on to show that if there is common knowledge of coherency (where i "knows" something if he assigns probability 1 to it), then i's type will determine not only his belief about j's type, but his belief about j's belief about his type, and so on. To state this formally, let $T'' \times T''$ denotes the set of types for which there is common knowledge of coherency.

Proposition 4 *There is a homeomorphism* $g : T'' \to \Delta(S \times T'')$*.*

So assuming common knowledge of coherency, it is possible to have a *universal type space* (the space $T'' \times T''$) for which each player's type will specifies *all* his beliefs. In this sense, it is possible to have a "closed" model of incomplete information.

As noted above, few applied economic models explicitly model higher order uncertainty. Rather, types are drawn from a small subset of the universal type sapce. There is, however, a recent literature that asks whether standard models that have only first-order uncertainty lead to predictions that are *robust* to perturbations of higher order beliefs (see, for instance, Rubinstein, 1989). Indeed, this is one way to view the global games analysis — as questioning the robustness of coordination equilibria that are not risk-dominant — even though the global games model we have considered retained the standard modeling approach.

Bergemann and Morris (2004) pursue the robustness line of inquiry in the context of mechanism design. They ask whether when a mechanism that implements some outcome in a standard model with first-order uncertainty will implement the same outcome if types are drawn from the universal type space constructed above. Their answer is that for a mechanism to be robust in this sense, it must implement the desired outcome as an *ex post* equilibrium, which in the case of private values is equivalent to *dominant strategy implementation.*

References

[1] Bergemann, D. and S. Morris (2004) "Robust Mechanism Design," Yale University Working Paper.

[2] Brandenburger, A. and E. Dekel (1993) "Hierarchies of Beliefs and Common Knowledge," *J. Econ. Theory*, 59, 189–198.

[3] Carlsson, H. and E. van Damme (1993) "Global Games and Equilibrium Selection," *Econometrica*, 61, 989-1018.

[4] Kajii, A. and S. Morris (1997) "The Robustness of Equilibria to Incomplete Information, *Econometrica*, 65, 1283–1309.

[5] Frankel, D., S. Morris and A. Pauzner (2002) "Equilibrium Selection in Global Games with Strategic Complementarities," *J. Econ. Theory.*

[6] Frankel, D. and A. Pauzner (2000) "Resolving Indeterminacy in Dynamic Games," *Quarterly J. Econ.*

[7] Harsanyi, J. (1968) "Games with Incomplete Information Played by 'Bayesian' Players, I–III," *Management Science*, 14, 159–182, 320–334, 486–502.

[8] Levin, J. (2001) "A Note on Global Equilibrium Selection in Overlapping Generations Games," Working Paper.

[9] Mertens, J. and S. Zamir (1985) "Formulation of Bayesian Analysis for Games with Incomplete Information," *International J. Game Theory*, 14, 1–29.

[10] Morris, S. and H. Shin (2002) "Global Games: Theory and Applications," *Advances in Economic Theory*, Proceedings of the Econometric Society Ninth World Congress.

[11] Morris, S. and H. Shin (2003) "Social Value of Public Information," *American Economic Review.*