INTERACTION GAMES: A Unified Analysis of Incomplete Information and Local Interaction

Stephen Morris

Sometimes models with very different motivation and outward appearance turn out to have close connections. Those connections may provide new insight into the original models. The TARK conferences have built on a leading example of this phenomenon: the striking connections between the models of interactive knowledge employed by philosophers, computer scientists, economists and game theorists. In these very different disciplines, analysts are concerned with what different agents (be they people or distributed processors) know about other agents' knowledge and how such interactive knowledge influences agents' ability to achieve co-ordinated behavior. Very different models - philosophers' epistemic logic, economists' state spaces and partitions, and computer scientists' models of distributed computing - turned out to be intimately related.

This paper also describes a new way of looking at the relation between different models. In particular, a precise connection between incomplete information games and local interaction games is described. Incomplete information games are the canonical way of modelling strategic environments in the presence of uncertainty about agents' preferences, agents' beliefs about other agents' beliefs about preferences, and so on. Local interaction games are a way of studying the way large populations of players interact strategically, without uncertainty, but when each player deals only with a small subset of the total population.

The simplest way to both introduce incomplete information and local interaction games, and describe the connection, is via an extended example. The incomplete information version of the example is a variation on the electronic mail game of Rubinstein [1989], which can be seen as a strategic version of the coordinated attack problem in the computer science literature. The local interaction version builds on an example of Ellison [1993].

Two players (*ROW* and *COL*) must choose action "Invest" (I) or action "Don't Invest" (D). Each player faces a cost 2 of investing. Each player realizes a gross return of 3 from the investment if both (1) the other player invests and (2) investment conditions are *favorable* for that player. Thus if investment conditions are favorable for both players, then payoffs are given by the following symmetric matrix:

Favorable for *ROW* Favorable for *COL*	I	D
I	1,1	-2,0
D	0,-2	0,0

This game has two strict Nash equilibria: both players invest and both players don't invest. On the other hand, if conditions are unfavorable for *ROW* (but favorable for player *COL*), payoffs are given by the following matrix:

Favorable for *ROW* Unfavorable for *COL*	I	D
I	-2,1	-2,0
D	0,-2	0,0

In this game, *ROW* has a dominant strategy to not invest, and thus the unique Nash equilibrium has both players not investing.

First consider the two players' ability to co-ordinate their behavior when there is a small amount of incomplete information about investment conditions. In particular, investment conditions are always favorable for *COL*, but not for *ROW*. *ROW* knows when investment conditions are favorable for him, but *COL* does not.

Specifically, suppose that *ROW* observes a signal $s_R \in \{0, ..., K-1\}$ which is drawn from a uniform distribution. Assume that investment conditions are favorable for *ROW* *unless* $s_R = 0$. *COL* observes a noisy version of *ROW*'s signal, $s_C \in \{0, ..., K-1\}$. In particular, assume that

$$s_C = \begin{cases} s_R, \text{ with probability } 1/2 \\ s_R - 1, \text{ with probability } 1/2 \end{cases},$$

with mod K arithmetic, so that $0 - 1 = K - 1$. Thus if $s_R = 0$, s_C is 0 or $K - 1$ with equal likelihood.

The above constitutes a description of an incomplete information game. We can summarize the game in the following diagram:

Type of *COL*

Type of *ROW*

	0	1	2		K-1	
0	$\times$	$\circ$	$\circ$	·	$\times$	U
1	$\times$	$\times$	$\circ$	·	$\circ$	F
2	$\circ$	$\times$	$\times$	·	$\circ$	F
	·	·	·	·	·	·
K-1	$\circ$	$\circ$	$\circ$	·	$\times$	F
	F	F	F	·	F	

Types of *ROW* are represented by rows, types of *COL* by columns. Boxes with a $\times$ correspond to type profiles which occur with positive probability; given the uniform prior assumption, each occurs with ex ante probability $\frac{1}{2K}$. Boxes with a $\circ$ correspond to type profiles that occur with zero ex ante probability. Payoffs are specified by the letter - F for favorable, U for unfavorable - at the end of the row/column corresponding to the type.

The unique equilibrium of this incomplete information game has each player never investing. To see why, observe first that type 0 of *ROW* will not invest in any equilibrium. But type 0 of *COL* attaches probability 1/2 to *ROW* being of type 0, and therefore not investing. But even if investment conditions are favorable, the best response of a player who believes that his opponent will invest with probability less than or equal to a half is not to invest. Thus type 0 of *COL* will not invest. But now consider type 1 of *ROW*. Although investment conditions are favorable, he attaches probability 1/2 to his opponent not investing; so he will not invest. This argument iterates to ensure that no one will invest.

This example illustrates the fact that, in order for investment to be an equilibrium outcome, it is not enough that investment conditions are favorable for both players with high probability; nor is it enough that everyone know that everyone know... up to an arbitrary number of levels... that investment conditions are favorable for both players.

Now consider a local interaction version of this story. Suppose that there are $2K$ players situated on a circle. Player k interacts with his two neighbours, $k-1$ and $k+1$. We use mod $2K$ arithmetic, so that player $2K$'s neighbours are $2K-1$ and 1. Conditions are favorable for all players except the player at location 1. It is common knowledge for whom investment conditions are favorable.

Each player must decide whether to invest or not. His payoff is the sum of his payoff from his two interactions with each of his two neighbours. A strategy profile specifies which players invest, and which do not. A strategy profile is an equilibrium strategy profile if each player's action is a best response given the behaviour of his two neighbours.

This local interaction game can be summarized by the following table:

	2	4	6		2K	
1	×	○	○	·	×	U
3	×	×	○	·	○	F
5	○	×	×	·	○	F
	·	·	·	·	·	·
2K-1	○	○	○	·	×	F
	F	F	F	·	F	

A cross (×) marks a pair of players who interact with each other. Thus, for example, player 3 interacts with players 2 and 4 and no other player.

The unique equilibrium of this game has all players never investing. The argument is as for the incomplete information game. We know that the player at location 1 will never invest. Consider the player at location 2. Since one of his neighbours is not investing, his best response is not to invest. Similarly, the player at location 3 does not invest, and the argument iterates to ensure the result.

The above table is constructed in such a way as to identify an exact relationship between the incomplete information game and the local interaction game. In particular, the odd numbered players in the local interaction game play the role of *ROW*'s types in the incomplete information game, while the even numbered players play the role of *COL*'s types.

The full version of this paper (Morris [1997b]) identifies more generally what lies behind this connection. Incomplete information and local interaction share a common structure. A type or player interacts with various subsets of the set of all types/players. A type/player's total payoff is additive in the payoffs from these various interactions. The full version of this paper describes a general class of interaction games and shows how incomplete information games and local interaction games can be understood as special cases. Techniques and results from the incomplete information literature are translated into this more general framework. As a by-product, it is possible to give a new, complete characterization of equilibria robust to incomplete information, in the sense of Kajii and Morris [1997], in many player binary action co-ordination games. Only equilibria that are robust in this sense [1] can spread contagiously and [2] are uninvadable under best response dynamics in a local interaction system. A companion paper, Morris [1997a], uses these techniques to characterize features of local interaction systems that allow contagion.

References

[1] Ellison, G. (1993). "Learning, Local Interaction, and Coordination," *Econometrica* 61, 1047-1071.

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

[3] Morris, S. (1997a). "Contagion," CARESS Working Paper #97-01, University of Pennsylvania.

[4] Morris, S. (1997b). "Interaction Games: A Unified Analysis of Incomplete Information, Local Interaction and Random Matching," Santa Fe Institute Working Paper #97-08-072 E.

[5] Rubinstein, A. (1989). "The Electronic Mail Game: Strategic Behavior under Almost Common Knowledge," *American Economic Review* 79, 385-391.