Observable Strategies

Karl Wärneryd*

August 5, 2008

Abstract

The informal idea of a game in which the players observe each other's strategies prior to play has been popular in economics as well as moral philosophy, but on closer inspection turns out to be logically inconsistent. On the other hand, a notion of limited observability of strategies, which may be implemented in practice through strategic delegation, is consistent. We show that this latter construction is associated with a Folk Theorem-like result for one-shot games, and discuss the implications for, among other things, agent-based electronic commerce. *Journal of Economic Literature* Classification Numbers: C70, C72.

---

[*] Department of Economics, Stockholm School of Economics, Box 6501, SE—113 83 Stockholm, Sweden, and CESifo. Email: `Karl.Warneryd@hhs.se`. I thank Ken Binmore, Fabrizio Germano, Holger M Müller, Rosemarie Nagel, and Ariel Rubinstein for helpful comments.

> Not only can one develop empathy to better understand the feelings of one's staff so as to better motivate them, one can also telepathically scan one's competitors' minds to learn their secrets.
>
> Tim Rifat, *Remote Viewing and Sensing for Managers: How to Use Military Psiops for a Competitive Edge*, 2003.

1 Introduction

Suppose that before playing, say, the Prisoners' Dilemma, a player could observe his opponent's strategy, and the opponent his. In his influential work on moral philosophy, Gauthier [13] suggests that in such a situation a rational player should choose to cooperate if the opponent is observed also to be a conditional cooperator, and defect otherwise.

Variants of this idea are not uncommon in discussions of rational behavior in moral philosophy. (See, e.g., Howard [14] or Danielson [5].) In economics, Frank [10, 11] argues at length that a form of transparency of human agents should be expected to be a factor in social interaction, since evolutionary forces would have favored the development of physical characteristics that reliably signal a person's disposition or strategy. Actual evidence of this being the case includes the fact, reported by Ekman [6], that it is difficult to lie without giving it away through facial expressions that are beyond conscious control.

Appealing as it may seem at first glance, this justification of conditional cooperation raises some fundamental questions. In this paper we shall focus on the following. For the players in a game to observe each other's *strategies*, as such, is, as we shall argue, a logical impossibility.

Nevertheless, a form of partial transparency or observability of strategies can be made logically consistent. We show that this is all that is needed in order to generate conditional cooperation in the Prisoners' Dilemma. Indeed, more generally a "Folk-Theorem"-like result holds for such games of partial observability of strategies.

While it is doubtful that the notion of (partially) observable strategies is of

		Player 2	
		c	d
Player 1	c	$2,2$	$0,3$
	d	$3,0$	$1,1$

Table 1: The Prisoners' Dilemma.

great practical interest if interpreted literally, the sentiment expressed in the epigraph to this paper notwithstanding, we also discuss how this idea may be implemented as strategic delegation.

2 Transparency

Consider the familiar Prisoners' Dilemma (PD) game of Table 1. This game has the property that playing action d is a strictly dominant strategy for each player, in that it alone yields a player his highest payoff no matter what the other player does, while both players would be better off if they both played c. The game is quite trivial from a game-theoretic point of view, of course, being a rare example of a game that has a unique solution in dominant strategies. Nevertheless it has troubled many moral philosophers and political scientists, who see in it a stylized version of the quintessential problem of social interaction—a conflict between individual rationality and the common good.

Gauthier [13] suggests that the prospects for voluntary cooperation in one-shot interactions may not be as bleak as all that. For, he argues, it would be in the interest of the truly rational individual to develop a *disposition* to cooperate, conditional on the opponent having the same disposition, and otherwise play d, and furthermore to make this disposition public. Such an individual Gauthier calls a *constrained maximizer* (CM).

Gauthier feels we should really consider the new game, given in Table 2, constructed from the example PD by having it played by individuals who can choose between the CM disposition, which makes its behavior contingent on

		Player 2	
		CM	SM
Player 1	CM	$2,2$	$1,1$
	SM	$1,1$	$1,1$

Table 2: Gauthier's disposition game.

the disposition of the opponent, and the old, *straightforward maximizer* (SM) disposition that always plays d.

This game has a Nash equilibrium where both players adopt the CM disposition, inducing cooperation. For if one player were to deviate, the constrained maximizer would see this, dispositions being assumed public, and retaliate.

This account of conditional cooperation begs numerous questions. We shall focus on the following one: Where are all the other possible dispositions in Gauthier's game?

For clearly there must be more ways of conditioning behavior on the opponent's disposition than just CM and SM. We can immediately think of two more, one which plays c regardless of whether the opponent is CM or SM, and one which plays d if the opponent is CM and c if the opponent is SM. This makes four possible dispositions so far. But then CM and SM are incompletely specified dispositions, since they do not specify what to do if the opponent has one of the two new dispositions. And so on.

Gauthier's dispositions are not *strategies* of the transparency game in the orthodox sense of *complete contingent plans of action* that specify what to do in every situation that could arise in the game. But before we can tell what the consequences of this omission of possible behaviors from the game are for the possibility of conditional cooperation, we must ask what the *complete* strategy set might look like.

Consider a normal form game G with player set $N := \{1, 2, \ldots, n\}$ and finite action sets A_i. Is it possible to extend this game into one where prior to taking actions in G, the players observe each other's strategies? That is, is the sentence

> "Prior to taking actions in G, the players observe each other's strategies."

meaningful?

The answer to this question is no. This is easiest to see in a symmetric 2-player game with common action set A. Assume, by way of contradiction, that such an extended game exists, and let I be a player's set of information sets, or information partition, of that game. A pure strategy for a player is a mapping from his information partition into A, so his set of pure strategies is $S := \{s | s: I \to A\}$. Assuming a player observes the strategy choice of his opponent, we must have $I = S$. Hence S is the set of all mappings from S itself into A. The pure strategy set of the hypothetical extended game is therefore self-referentially defined by the equation

$$S = \{s | s: S \to A\}. \tag{1}$$

Does such a fixpoint set S exist? In the trivial case where A has a single element a, it does. Then the unique solution to equation (1) is the set $S = \{s\}$ where $s(s) = a$. In general, however, there is no solution, as can easily be proved using a Cantorian diagonalization argument.

Proposition 1 *Suppose we have $|A| \geq 2$. Then there is no set S satisfying the equation $S = \{s | s: S \to A\}$.*

Proof. Suppose there was a fixpoint set S. Consider a new mapping $s': S \to A$ such that $s'(s) \neq s(s)$ for all $s \in S$. Since A has at least two elements, this construction is always possible. The mapping s' does not belong to S, since it differs from every $s \in S$ at at least one point. So we have a contradiction. Therefore S cannot be a solution to (1). $\square$

That is, the game suggested by Gauthier has no strategy set.

By imposing restrictions on the allowable mappings one can, however, construct valid fixpoint sets. For instance, the equation

$$S = \{s | s \text{ is a constant function from } S \text{ into } A\}$$

has a fixpoint set with $|A|$ elements. McAfee [15], Binmore [2], Anderlini [1], and Canning [4] suggest that the question of the limits of rationality may be investigated by studying games played by Turing machines that input each other's descriptions before play. In this case we are dealing with a set of decision rules S such that

$$S = \{s | s \text{ is a Turing-computable function from } S \text{ into } A\}.$$

Such a set exists because there are only countably many Turing machines.

Gauthier's game with only two dispositions cannot be rationalized in such a fashion.

3 Semi-Public Strategies: An Example

In this section, we show that it is possible to construct a logically consistent model of a game where the players observe *part* of each other's strategies prior to taking action. Further, we exhibit an example that shows that this is all that is needed to generate rational conditional cooperation in dilemma situations.

Consider a symmetric, two-player normal form game G with common finite action set A, extended in the following manner. Before simultaneously taking action in G, the players each privately observe a *datum* consisting of a finite string of symbols from a finite alphabet $\mathscr{A}$. Suppose each datum is of length $k \geq 1$. Let $\mathscr{L}_k$ be the set of strings of length k of symbols from $\mathscr{A}$. Then, since the choice of action may be made contingent on the observed datum, a *pure strategy* for a player in the extended game is a mapping from $\mathscr{L}_k$ into A. That is, the common strategy set of the extended game is $S := \{s | s \colon \mathscr{L}_k \to A\}$. Since $\mathscr{L}_k$ and A are finite, S is well-defined and has $|A|^{|\mathscr{L}_k|}$ elements.

Now suppose we endow $\mathscr{L}_k$ with some arbitrary order. Then each element in S may be written as a string of symbols representing actions in A, such that the first position of the string representation of a strategy is a symbol representing the action taken when the datum string is the least element of $\mathscr{L}_k$, and so on.

		Player 2			
		$\underline{c}c$	$\underline{c}d$	$\underline{d}c$	$\underline{d}d$
Player 1	$\underline{c}c$	$2,2$	$2,2$	$0,3$	$0,3$
	$\underline{c}d$	$2,2$	$2,2$	$1,1$	$1,1$
	$\underline{d}c$	$3,0$	$1,1$	$2,2$	$0,3$
	$\underline{d}d$	$3,0$	$1,1$	$3,0$	$1,1$

Table 3: The Prisoners' Dilemma with semi-public strategies.

Our next step is to assume that $\mathscr{A}$ is the symbol alphabet representing actions in A, and that the datum string observed by a player is a part of the string representation of his opponent's strategy. Observe that, by construction, a datum is always a proper substring of the string representation of a strategy. If the datum length is k, and A (and therefore $\mathscr{A}$) has $m \geq 2$ elements, then there are m^k different strings in $\mathscr{L}_k$, so the symbolic representation of a strategy is of length $m^k > k$. We also note that, in one sense, the relative information provided by a datum is decreasing in its length, in that we have for the ratio of datum length to strategy length that $\lim_{k\to\infty} k/m^k = 0$.

We now have a consistent model of a game in which the players may be said to observe part of each other's strategies prior to play. To see the implications of this construction, consider the following example. Let G be the game of Table 1, let $\mathscr{A} = \{\text{“}c\text{”}, \text{“}d\text{”}\}$, and suppose we have $k = 1$. A datum is now a single symbol, so a strategy may be written as a string xy, where we take $x \in \mathscr{A}$ to be the symbol representing the action chosen when the datum is “c”, and $y \in \mathscr{A}$ that chosen when the datum is “d”. Finally, we identify a player's datum with the first symbol of the string representation of his opponent's strategy. Table 3 shows the normal form of the resulting extended game, with the public parts of the strategies underscored.

The extended game has two pure-strategy equilibria, (cd, cd) and (dd, dd), which correspond to the equilibria of Gauthier's game. We thus conclude that full transparency of decision procedures, even if this notion could be made log-

		Player 2			
		$c\underline{c}$	$c\underline{d}$	$d\underline{c}$	$d\underline{d}$
Player 1	$c\underline{c}$	2,2	2,2	0,3	0,3
	$c\underline{d}$	2,2	1,1	2,2	1,1
	$d\underline{c}$	3,0	2,2	1,1	0,3
	$d\underline{d}$	3,0	1,1	3,0	1,1

Table 4: The Prisoners' Dilemma with different semi-public strategies.

ically consistent, is not necessary for conditional cooperation to be a stable outcome. What matters is that a player knows what the opponent will do if the first player cooperates. Below, we shall show that, in fact, any outcome that is individually rational in a sense to be made precise is an equilibrium outcome of some game extended with semi-public strategies.

Which part of the strategy is public is crucial for the possibility of cooperation, however. Note that the example in Table 3 illustrates the situation when the information revealed is what a player would do if his opponent turned out to be a conditional cooperator. The possibility of cooperation arises here because the intent to cooperate will be common knowledge among two conditional cooperators.

Consider now instead the extended game that results if the information revealed is the *second* entry in a player's strategy. This case is illustrated in Table 4.

In this game, the unique equilibrium is for both players to choose the strategy dd. The prospects for cooperation vanish entirely. It should be noted that, in particular, the strategy cd of this game is *not* a conditional cooperator in this game. In fact, this game does not have a strategy corresponding to conditional cooperation.

4 Semi-Public Strategies: Formalities

We now generalize the notion of semi-public strategies developed in the example.

Let G be a normal form game with pure strategy, or *action*, set A_i for player $i \in N := \{1, 2, \ldots, n\}$. We assume $|A_i| \geq 2$ for all i. Write $A := \times_i A_i$. For $a \in A$, $u_i(a) \in \mathbb{R}$ is player i's payoff.

We construct from G the extended game $G^\star := \langle G, (\Sigma_i), (\varphi_i), (P_i), (\prec_i)\rangle$, where

- Σ_i is an alphabet of symbols such that $|\Sigma_i| = |A_i|$,
- $\varphi_i : A_i \to \Sigma_i$ is a one-to-one mapping,
- P_i is a subset of the natural numbers of cardinality k_i such that $\max P_i \leq \prod_{j \neq i} |A_j|^{k_j}$, and
- $\prec_i$ is an order on $\bar{L}_{\sim i} := \times_{j \neq i} L_j$, where L_i is the set of all strings of length k_i of symbols from Σ_i.

We give $G^\star$ the following interpretation. In $G^\star$, first each player i observes a *datum* $\mathbf{m}_{\sim i} \in \bar{L}_{\sim i}$. Subsequently, the players play G.

Observe that $G^\star$ is not yet strictly speaking a game, since we have not specified where the data come from, i.e., if they are the results of choices by Nature or the players themselves. Still, a *decision rule* of $G^\star$ can be usefully defined as a mapping $s_i : \bar{L}_{\sim i} \to A_i$. Now note that s_i can be written as a string of symbols $\varphi_i(s_i(\mathbf{m}^1))\varphi_i(s_i(\mathbf{m}^2))\cdots$, where $\mathbf{m}^1$ is the least element of $\bar{L}_{\sim i}$ according to $\prec_i$, $\mathbf{m}^2$ is the successor of $\mathbf{m}^1$, and so on. Abusing notation slightly, write $\varphi_i(s_i)$ for the string representation of s_i. Given $k_1, \ldots, k_n$, the length of $\varphi_i(s_i)$ is $|\bar{L}_{\sim i}| = \prod_{j \neq i} |A_j|^{k_j} > k_i$. Let $\varphi_i|_{P_i}(s_i)$ be the string resulting from picking out only the positions indicated by P_i. Our final step is to assume that

$$\mathbf{m}_{\sim i} = (\varphi_1|_{P_1}(s_1), \ldots, \varphi_{i-1}|_{P_{i-1}}(s_{i-1}), \varphi_{i+1}|_{P_{i+1}}(s_{i+1}), \ldots, \varphi_n|_{P_n}(s_n)).$$

The payoffs in $G^\star$ are those resulting from the play of G. Again abusing notation slightly, extend u_i so that $u_i(s)$ is the payoff resulting from the actions in G induced by s.

To recapitulate, the idea here is that a pure strategy of a normal form game associates a unique action with every information set. Thus a strategy may be written as a list of actions. We introduce the notion that each player observes parts of the lists chosen by the other players.

Given this set-up, we may simplify analysis considerably by noting the following. Since given the φ_i every $\mathbf{m}_{\sim i}$ corresponds to a unique element of $\times_{j\neq i} A_j^{k_j}$, we may write the strategy set of player i directly as

$$S_i := \{s_i | s_i : \times_{j\neq i} A_j^{k_j} \to A_i\}.$$

This set is well defined and has $|S_i| = |A_i|^{|\times_{j\neq i} A_j^{k_j}|}$ elements.

Furthermore, given the P_i and k_i, we may associate with player i the set $\bar{A}_i \subset \times_{j\neq i} A_j^{k_j}$ such that $|\bar{A}_i| = k_i$. $\bar{A}_i$ is the set of possible observations about the other players' strategies to which player i's responses are public.

We define the *play* of player i at $s \in S$ in $G^\star$ to be

$$\phi_i(s) = s_i(s_1(\bar{a}_1^1), \dots, s_1(\bar{a}_1^{k_1}), \dots, s_n(\bar{a}_n^1), \dots, s_n(\bar{a}_n^{k_n}))$$

where $\bar{a}_j^\ell \in \bar{A}_j$.

5 A "Folk Theorem"

We can now characterize the equilibria of games with semi-public strategies. Define

$$\bar{u}_i := \max_{a_i \in A_i} \min_{a_{\sim i} \in A_{\sim i}} u_i(a_i, a_{\sim i}),$$

the *maximin* payoff of player i in G.

Proposition 2 *Let $a \in A$ be an action profile of G. If and only if we have $u_i(a) \geq \bar{u}_i$ for all i, there is an extended game $G^\star$ such that there is a Nash equilibrium of $G^\star$ that induces a.*

Proof. Sufficiency: Let $a := (a_1, a_2, \dots, a_n)$ be such that $u_i(a) \geq \bar{u}_i$ for all i, let $k_i = 1$ for all i, and let

$$\phi_i(s) = s_i(s_1(a_{\sim 1}), \dots, s_{i-1}(a_{\sim(i-1)}), s_{i+1}(a_{\sim(i+1)}), \dots, s_n(a_{\sim n})).$$

Define

$$\hat{a}_{\sim i}(a'_i) := \arg\min_{x_{\sim i}\in A_{\sim i}} u_i(a'_i, x_{\sim i}).$$

Consider s such that $s_i(a_{\sim i}) = a_i$ and $s_i(a_1,\ldots,a'_j,\ldots,a_n) \in (\hat{a}_{\sim j}(a'_j))_i$ for all $j \neq i$ and all $a'_j \neq a_j$. Since we have $u_i(s) = u_i(a) \geq \bar{u}_i \geq u_i(s_{\sim i}, s'_i) = u_i(\hat{a}_{\sim i}(a'_i))$ for all s'_i such that $s'_i(a_{\sim i}) = a'_i \neq a_i$, s is an equilibrium that induces a.

Necessity: Suppose a is such that $u_i(a) < \bar{u}_i$ for some i. Let s be such that $\phi_j(s) = a_j$ for all j. Consider the strategy $s'_i \neq s_i$ such that

$$\phi_i(s'_i, s_{\sim i}) \in \arg\max_{a'_i\in A_i} \min_{a'_{\sim i}\in A_{\sim i}} u_i(a'_i, a'_{\sim i}).$$

We must have $u_i(s'_i, s_{\sim i}) \geq \bar{u}_i > u_i(s)$, so s'_i is a strictly better reply to $s_{\sim i}$ than s_i. Hence s cannot be an equilibrium. □

This "Folk Theorem" is somewhat analogous to those of repeated game theory. In a repeated game, a particular equilibrium may be sustained by the threat of future reversion to something worse if some player deviates in the current period. The current setting essentially collapses the repetition into a one-shot game, since it allows a player's intention to deviate to be punished immediately.

The intuition for how this works is as follows. If we have an action profile a that gives each player at least his maximin payoff, then an extended game can be constructed where each player reveals only a single bit of information about his strategy. This bit, specifically, is the action the player will take if he observes $a_{\sim i}$ from the other players. The profile a can be supported in equilibrium because strategies can be constructed that respond with actions that minimize player i's payoff if he reveals an action other than a_i. Hence the best payoff a deviating player could achieve is his maximin payoff. Because the strategies of the punishing players can be made contingent on the deviation action, but the deviation cannot be made contingent on the planned punishment, which is not observed in this construction, a deviating player can be held down to his maximin payoff. This contrasts, of course, with the folk theorems of infinitely repeated games (e.g., Fudenberg and Maskin [12]), in which a deviating player can only be held down to his minmax stage-game payoff. In that setting, actions during the punishment phase are selected independently, and cannot be made contingent on each other, so that the player being punished can always play a best reply to the punishment actions of other players.

As we have already seen by means of an example, it is of course also always possible to construct extended games that do *not* induce the desired action profile in any equilibrium.

That an equilibrium of any game extended with semi-public strategies must necessarily give each player at least his maximin payoff follows because regardless of the other players do, a player can always guarantee himself the maximin payoff.

6 An Application: Strategic Delegation

The idea of semi-public strategies does not make much sense if taken literally, since it seems to require mind-reading skills of the game-players. An interpretation that is perhaps more appealing is to think of the players as *agents* who play the underlying game on behalf of some background *principals.* The notion of such *strategic delegation* is discussed in, e.g., Fershtman and Judd [7], Fershtman, Judd, and Kalai [8], Fershtman and Kalai [9], and Caillaud, Jullien, and Picard [3].

On this reading, we may think of the agents' strategies as sets of direct instructions received from the respective principals. These instructions bind the agent to take certain actions conditional on what is known about the contracts of the other parties. To visualize this, think of the instructions as being written down in some commonly acknowledged language on sheets of paper. The sheets of paper are then partially slid into envelopes that are kept in open view. The parts of the contracts that stick out of the envelopes correspond to the public elements of the agents' strategies.

The proof of Proposition 2 now directly suggests how any outcome that gives each principal more than his maximin payoff may be achieved through delegation. Suppose there is some such action profile a that the principals wish to achieve, but which they could not implement by contracting directly with each other. Set $k_i = 1$ for all i, and let the public part of the contract signed by party i deal with the action taken if the public parts of the contracts of the other parties spell out $a_{\sim i}$. Let the instruction in this case be a_i, and an element of $(\arg\min_{x_{\sim j} \in A_{\sim j}} u_j(a'_j, x_{\sim j}))_i$ if some party j specifies something other than a_j. A set of such contracts will be an equilibrium of the delegation game that induces play of a. Since the contracts as such are observable, this delegation equilibrium is also robust against renegotiation.

A similar “Folk Theorem”-like result for delegation games is discussed by Fershtman, Judd, and Kalai [8], but the mechanism is slightly different. Fershtman *et al* assume that a contract specifies a remuneration level for the agent as a function of the payoff of the principal, and that contracts are observable in the sense that an agent may condition his action choice on the contracts signed.

References

[1] Luca Anderlini. Some notes on Church’s thesis and the theory of games. *Theory and Decision*, 29:19–52, 1990.

[2] Ken Binmore. Modeling rational players: Part I. *Economics and Philosophy*, 3:179–214, 1987.

[3] B Caillaud, B Jullien, and P Picard. Competing vertical structures: Precommitment and renegotiation. *Econometrica*, 63:621–646, 1995.

[4] David Canning. Rationality, computability, and Nash equilibrium. *Econometrica*, 60:877–888, 1992.

[5] Peter Danielson. *Artificial Morality: Virtuous Robots for Virtual Games*. Routledge, London, 1992.

[6] Paul Ekman. *Telling Lies*. W W Norton & Company, New York, 1985.

[7] Chaim Fershtman and Kenneth L Judd. Equilibrium incentives in oligopoly. *American Economic Review*, 77:927–940, 1987.

[8] Chaim Fershtman, Kenneth L Judd, and Ehud Kalai. Observable contracts: Strategic delegation and cooperation. *International Economic Review*, 32:551–559, 1991.

[9] Chaim Fershtman and Ehud Kalai. Unobserved delegation. Working paper, Northwestern University, 1994.

[10] Robert H Frank. If *homo economicus* could choose his own utility function, would he want one with a conscience? *American Economic Review*, 77:593–604, 1987.

[11] Robert H Frank. *Passions Within Reason: The Strategic Role of the Emotions*. W W Norton & Company, New York, 1988.

[12] Drew Fudenberg and Eric Maskin. The folk theorem in repeated games with discounting or with incomplete information. *Econometrica*, 54:533–554, 1986.

[13] David Gauthier. *Morals By Agreement*. Oxford University Press, Oxford, 1986.

[14] Nigel Howard. *Paradoxes of Rationality: Theory of Metagames and Political Behavior*. The MIT Press, Cambridge, MA, 1971.

[15] R Preston McAfee. Effective computability in economic decisions. Working paper, University of Western Ontario, 1984.