

Observable Actions¹

Ryosuke Ishii²

Assistant Professor, Faculty of Business, Aichi Shukutoku University

23 Sakuragaoka, Chikusa-ku, Nagoya, Aichi 464-8671, Japan

November 12, 2012

¹I am grateful to Akira Okada, Haruo Imai, Tadashi Sekiguchi, and Kenichi Amaya for their detailed comments. I also thank seminar participants at Fukuoka University, Otaru University of Commerce, Kwansei Gakuin University, Nagoya University, Konan University, Tohoku University, and Waseda University for their constructive comments and suggestions. All errors in this paper remain mine.

²Email: rishii@asu.aasa.ac.jp Phone: +81 52 781 1151 (Switchboard)

Abstract

We consider a game with “meta-players” who observe each other’s actions before actual play. This observability exerts an effect similar to repeated games without discounting. This analysis is novel in that it makes “mimic strategies” possible; meta-players are allowed to take the same actions as opponents would take. Such mimic strategies have been excluded from strategy sets as a cause of an indeterminacy problem inherent in meta-game settings in the existing literature. We resolve the problem by introducing “beliefs” about actions that opponents are taking. The game has Nash equilibria with any individually rational payoff profiles. In addition, the outcomes that satisfy a modified version of evolutionary stability lead to Pareto efficiency in coordination games.

Keywords: Metagame; Folk theorem; Evolutionarily stable set; Equilibrium selection

JEL classification: C70

1 Introduction

In the Prisoners' Dilemma game, playing "confession" is the unique strictly dominant strategy although the outcome (confession, confession) is not Pareto-efficient. How to solve such dilemma-like situations, one of which is the trigger strategies in infinite repeated games, has been discussed. In this paper, we present a different approach to this problem.

The idea behind our argument is of ancient vintage. Gauthier (1986) suggests that truly rational players can develop *dispositions*, and that voluntary cooperation in one-shot interactions could be achieved. Suppose a disposition to take "silence" if the opponent has the same disposition, and take "confession" if different. We call this disposition a *constrained maximizer*. Another disposition is to take "confession" always. We call it a *straightforward maximizer*. Such dispositions are observable before taking actions. Players choose not between actions "silence" and "confession," but between the above two dispositions.

There is a Nash equilibrium in this new game with a pair of the constrained maximizer dispositions, and induces the Pareto-efficient outcome. However, this formularization is open to question: Where are all the other possible dispositions in Gauthier's (1986) game? Other conditional behaviors on opponent's disposition could be possible than the above two dispositions. For example, one can play "silence" always.

Once we introduce the new disposition, the specification of the above two dispositions are incomplete. There would be multiple ways to specify what to do if the opponent has the new disposition, and so on.

Howard (1971) defines “metagame” as formal game theory. Let us consider 1-metagame first. “1-” means player 1’s conditional strategy. Player 1 does not play simple action “confession” or “silence,” but instead chooses a mapping from player 2’s actions to self actions. Player 1’s possible strategies are playing “confession” always, playing “silence” always, playing the same action as player 2, and playing the opposite to player 2.

We can make the problem even more complicated. Let us assume that player 2 can make a choice based on player 1’s choice in 1-metagame, here referred to as 21-metagame. Then player 2’s strategies are mappings from four player 1’s strategies to the set of player 2’s actions (“confession” and “silence”).

In this 21-metagame, there are three pure strategy Nash equilibria, two of which correspond to a Pareto-efficient outcome (silence, silence), and one of which corresponds to a Pareto-dominated outcome (confession, confession). Howard (1971) concludes that the Pareto-efficient outcome stands to reason since no one prefers Pareto-dominated outcomes.

There are, however, serious doubts. At first, Howard (1971) regards

it as trivial to narrow down the candidates of a solution to Pareto-efficient ones. This is not necessarily the case on equilibrium selection, for example, in coordination games.

More seriously, we may get into questions about how multilayer the players' inferences are. Is it sufficient to suppose the "21" layer? Or, is there any metagame with symmetric strategies?

Unfortunately, we cannot construct a game in which players observe each other's strategy before playing games, from a strictly logical perspective. (See, for example, Wärneryd (2010)).

There are two approaches to avoid this problem. First, Kalai et al (2010) restrict strategies to play, and prove a folk-theorem-like result. However, the above problem in Gauthier (1986) occurs once again. That is, we cannot justify a strategy set that excludes some strategies. Kalai et al (2010) write in their example of price competition "This formulation disallows vague ads, like 'I will undercut opponents' prices by \$50,' which fail to specify a response price to an identical competitor's ad." This restriction of the strategy set might be too restrictive. Although there is no consistent action pair in the case that both players will undercut each other, consistent outcomes exist when one will undercut the other and when the other will overcut (setting aside whether the player is willing to "overcut"). When we interpret the whole game as a delegation

agreement, the undercut strategy seems natural. It may distort the game to suppose that players can never adopt the undercut strategies. To the author's best knowledge, this paper is the first study that provides a model with "undercut strategies" without causing the indeterminacy problem.

The second approach is to observe not strategies themselves but a coarse partition of each other's strategy set as in Wärneryd (2010). For example, we can consider a game in which before playing a Prisoners' Dilemma game, the players observe to which "class" between the following two below the opponent's strategy belongs:

class 1 Silence when the opponent's strategy belongs to class 1 and confession when the opponent's strategy belongs to class 2.

class 2 All the other strategies.

Certainly, other classifications can be created by changing the partition of the strategy set.

Wärneryd (2010) shows that for any underlying 2-player, finite, normal-form game there is a game extended with such coarsely observable strategies that has equilibria with payoffs arbitrarily close to any feasible, individually rational payoff profile. However, there is no discussion as to whether the classifications are reasonable. In this example, when a player observes class 1, she comes to know the opponent's strategy before the

actual play. Otherwise, she cannot know for certain which strategy the opponent would play, since the strategy belongs to class 2 with multiple strategies. The classification discriminates against actions or strategies.

We will construct a setting of observable actions. It is well defined that the players observe one another's actions (which the opponents are actually playing). In our setting, the "undercut strategies" are (potentially) possible for players, and there is no distinction among strategies.

Since observability of actions allows players to send signals to one another, it is natural to compare our model and pre-play communication games. In communication equilibria all players send costless signals to one another, whereas in our model, players can send only commitment signals. Although pre-play communication extension in games with multiple equilibria has been successful in leading to Pareto efficiency via a "secret handshake" to some extent (e.g., Wärneryd (1991)), there is still an inefficient equilibrium that is evolutionarily stable. Our model replaces cheap talk with observability, and resolves the problem. That is, only Pareto-efficient outcomes are selected among multiple equilibria.

2 Model

2.1 Underlying Game

We consider a two-person normal form game $G = \langle N = \{1, 2\}, (A_1, A_2), (u_1, u_2) \rangle$.

N is a set of players. A_i is an action set for player $i \in N$. Assume $|A_i| \geq 2$

for all i . We define a set $A = (A_1, A_2)$ of action profiles. $u_i : A \rightarrow \mathbb{R}$ is a

utility function for player i . We use the standard convention where for

every player i , player $-i$ denotes the other player.

Example 1 Consider a 2×2 prisoners' dilemma:

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

The action set for each player is $\{c, d\}$.

2.2 Meta Game

Now we introduce meta strategies, assuming that each player observes

the opponent's action, and takes various actions depending on the op-

ponent's actions. There may be multiple action profiles realized when a

meta strategy simply means an action plan; that is, the mapping from

an opponent's action set to her action set. In a symmetric game, for example, the mimic plan, "I would take the same action as the opponent would take" is such a mapping. When both players employ mimic plans, all symmetric action profiles can be realized. We need a kind of "beliefs" (see below) to determine the actual opponent's action.

Example 2 We consider the same 2×2 prisoners' dilemma as mentioned above. A strategy can be represented as a string xy , where x (resp. y) is the action chosen when the opponent's action is c (d). We identify a player's belief about the opponent's action with the capital letter, which has to be the action she plays actually. For example, strategy Cd means to take action c if the opponent's action is c , to take d if the opponent's action is d , and to have a belief that the opponent would take c (because the first letter of the string is capitalized).

Formally, we define a meta game $G^* = \langle N, \Sigma = (\Sigma_1, \Sigma_2), (f_1, f_2) \rangle$. $\sigma_i = (s_i, b_i) \in \Sigma_i = (S_i, B_i)$ is a strategy for player i , where $\sigma = (\sigma_1, \sigma_2) \in (\Sigma_1, \Sigma_2) = \Sigma$. $s_i : A_{-i} \rightarrow A_i$ is an action plan for player i , where $s = (s_1, s_2) \in (S_1, S_2) = S$. $b_i \in B_i = A_{-i}$ is a belief for player i , where $b = (b_1, b_2) \in (B_1, B_2) = B$.

Example 3 Strategy profile (Cd, Cd) is well defined in the sense that the belief of each player is consistent with the opponent's action, which

we call a feasible strategy profile. The realized payoff profile is $(2, 2)$. Contrarily, (Cd, Dc) is not feasible. There is no belief profile that would allow this mapping pair to be feasible.

Definition 1 Strategy profile σ is feasible in G^* if $s_{-i}(b_{-i}) = b_i$ for all i .

In the case that σ is feasible, $f_i(\sigma) = u_i(s_i(b_i), b_i)$ holds for all $i \in N$.

Searching for feasible strategy profiles in meta games and for Nash equilibria in underlying games have much in common. When we interpret a Nash equilibrium in an underlying game, we usually suppose that players have conjectures about others' actions. Players choose actions as best responses to their conjectures. When the play begins, the conjectures are consistent with the realized action profile.

Definition 2 $\sigma'_i = (s'_i, b'_i)$ is a feasible deviation from feasible strategy profile $\sigma = (s, b)$ if there exists a (unique) belief $b''_{-i} \in B_{-i}$ such that ex-post strategy profile $\sigma|\sigma'_i = ((s'_i, s_{-i}), (b'_i, b''_{-i})) \in \Sigma$ is feasible.

The description of deviations in our meta game is akin to the one in repeated games. Action plan s'_i and action profile $(s'_i(b'_i), b'_i)$ correspond to a strategy and action profile flow, respectively, on the equilibrium path in repeated games. When a player deviates and changes her strat-

egy, the action profiles on the equilibrium path may change despite the opponent's strategy remaining unchanged. Thus, even a "unilateral" deviation can cause a change in the opponent's realized action.

Definition 3 Strategy profile σ is a Nash equilibrium if for any i and feasible deviation σ'_i ,

$$f_i(\sigma|\sigma'_i) \leq f_i(\sigma)$$

holds.

Example 4 Player 1 has four feasible deviations from strategy profile (Cd, Cd) . Here is a list of the ex-post strategy profiles corresponding to all player 1's feasible deviations, and the payoff profiles:

(Cd, Cd) (itself!)	$(2, 2)$
(Cc, Cd)	$(2, 2)$
(cD, cD)	$(1, 1)$
(dD, cD)	$(1, 1)$

The same holds for player 2. Therefore, (Cd, Cd) is a Nash equilibrium.

Consider meta strategy (s_i, b_i) in which $s_i(b'_i) = a_i$ for all b'_i as a meta extension of action a_i in the underlying game.

Lemma 1 A strategy profile in which all strategies are meta extensions

of actions consisting of a pure strategy Nash equilibrium in underlying game G is a Nash equilibrium in meta game G^* .

Using a (payoff) matrix of an $n \times m$ underlying game may be easier to understand. Suppose that player 2 plays meta strategy (s_2, a_1^1) as an extension of action a_2^1 , and that player 1 plays (s_1, a_2^1) in which $s_1(a_2^1) = a_1^1$. The action profile realized is (a_1^1, a_2^1) , which corresponds to mark \heartsuit in the payoff matrix below. The beliefs in feasible deviations by player 1 have to be a_2^1 . In the matrix, the action profiles realized by the feasible deviations by player 1 line up on the first column, which corresponds to mark \spadesuit (or \heartsuit).

		Player 2			
		a_2^1	a_2^2	\cdots	a_2^m
	a_1^1	\heartsuit	.	\cdots	.
Player 1	a_1^2	\spadesuit	.	\cdots	.
	\vdots	\vdots	\vdots	\ddots	\vdots
	a_1^n	\spadesuit	.	\cdots	.

In contrast, if player 2 plays meta strategy (s'_2, a_1^1) in which $s'_2(a_1^1) = a_2^1$, $s'_2(a_1^2) = a_2^m, \dots$, and $s'_2(a_1^n) = a_2^2$, the corresponding action profiles

realized by the feasible deviations by player 1 are as follows:

		Player 2			
		a_2^1	a_2^2	\cdots	a_2^m
	a_1^1	♥	·	\cdots	·
Player 1	a_1^2	·	·	\cdots	♠
	\vdots	\vdots	\vdots	\ddots	\vdots
	a_1^n	·	♠	\cdots	·

When we consider feasible deviations in a meta game, we must focus not only to the column in the matrix in the underlying game, but also to all cells of the matrix.

3 Folk Theorem

We can now characterize the equilibria of meta games. Define

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

as the maximin payoff of player i in underlying game G .

Proposition 1 Let $a \in A$ be an action profile of underlying game G .

If and only if we have $u_i(a) \geq \bar{u}_i$ for all i , there is a Nash equilibrium of meta game G^* that induces a .

Proof. Sufficiency: Define

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

Consider $\sigma = (s, b)$ such that $s_i(a_{-i}) = a_i$, that $s_{-i}(a'_i) \in \hat{a}_{-i}(a'_i)$ for all $a'_i \neq a_i$, and that $b_i = a_{-i}$ for all i . Since we have $f_i(\sigma) = u_i(a) \geq \bar{u}_i \geq u_i(\hat{a}_{-i}(a'_i)) = f_i(\sigma | \sigma'_i)$ for all feasible deviations $\sigma'_i = (s'_i, b'_i)$ such that $s'_i(a_{-i}) = a'_i \neq a_i$, σ is a Nash equilibrium that induces a .

Necessity: Suppose a such that $u_i < \bar{u}_i$ for some i . Let $\sigma = (s, b)$ such that $s_j(b_j) = a_j$ for all j . There exists feasible deviation $\sigma'_i = (s'_i, b'_i) \neq \sigma_i$ such that

$$s'_i(b'_i) \in \arg \max_{a'_i \in A_i} u(a'_i, s_{-i}(a'_i)),$$

so that σ'_i is a strictly better reply to σ_{-i} than σ_i . Hence σ cannot be an equilibrium. ■

4 Equilibrium Selection

In this section we focus on a finite and symmetric two-person game with a unique Pareto dominant payoff profile. We extend the concept of an evolutionarily stable set to our meta game. The basic idea follows:

There are many agents who play identical strategies in a meta game.

If a small number of agents mutate and play other (identical) strategies, incumbents are not defeated. If there is a strategy by which mutants keep the same performance as the incumbents, it is included in the evolutionarily stable set.

Defining evolutionary stability in our meta game causes a delicate problem relating to beliefs. Thus, we define stability not by single-population but by two-population.

Definition 4 A set X of strategy profiles is an evolutionarily stable set if for all $\sigma \in X$, for all $i, j = 1, 2$ ($i \neq j$), for all feasible deviations τ_i from σ , and for all feasible deviations τ_j from $\sigma|\tau_i$ such that the action plan of τ_j is the same as that of τ_i ,

1. $f_i(\sigma|\tau_i) < f_i(\sigma)$ or
2. $f_i(\sigma|\tau_i) = f_i(\sigma)$ and $f_j(\sigma|\tau_i|\tau_j) < f_j(\sigma|\tau_i)$ or
3. $f_i(\sigma|\tau_i) = f_i(\sigma)$ and $f_j(\sigma|\tau_i|\tau_j) = f_j(\sigma|\tau_i)$ and $\sigma|\tau_i|\tau_j \in X$.

It is straightforward to check that two evolutionarily stable sets coincide or are disjoint using the same type of reasoning as with the usual concept of evolutionarily stable sets.

The Pareto-efficient outcome in such a meta game is clearly that both

players receive a Pareto-efficient payoff α in the underlying game. Let

$$X^\# = \{\sigma \in \Sigma : f_1(\sigma) = f_2(\sigma) = \alpha\}.$$

We denote $c \in A_1 = A_2$ as an action that induces α . In general, this set is not a singleton.

Example 5 We consider a 2×2 symmetric coordination game:

		Player 2	
		c	d
Player 1	c	2, 2	0, 0
	d	0, 0	1, 1

Nash equilibrium (dD, dD) is not in the evolutionarily stable set of this meta game. Suppose that it is evolutionarily stable. Then there are feasible deviations $\tau_1 = \tau_2 = cD$, and (cD, cD) must be in the evolutionarily stable set. However, this is not a Nash equilibrium since Cd is a feasible deviation for both players and strictly improves the payoffs, which leads to a contradiction.

Proposition 2 $X^\#$ is the unique evolutionarily stable set in G^* .

Proof. We prove first that $X^\#$ is an evolutionarily stable set. Suppose that $\sigma \in X^\#$, and let τ_i and τ_j be feasible deviations from σ and $\sigma|\tau_i$

respectively. We proceed to show that one of the conditions 1.-3. in the above definition is met. σ is a Nash equilibrium since players could not obtain higher payoffs than α : $f_i(\sigma|\tau_i) \leq f_i(\sigma)$. If the inequality is strict, condition 1. is met. In the case of equality, we have $f_i(\sigma|\tau_i) = f_i(\sigma) = \alpha$. However, since α is the Pareto-efficient payoff, τ_i must take only an action that makes σ_j react with c , and τ_i must always reply to the opponent's action from σ by taking c . Hence, $f_j(\sigma|\tau_i) = \alpha$, and thus $f_j(\sigma|\tau_i|\tau_j) \leq f_j(\sigma|\tau_i)$ by Pareto dominance. If the inequality is strict, condition 2. is met. If equality holds, then $f_j(\sigma|\tau_i|\tau_j) = \alpha$, which implies, as before, that τ_j reacts with c to $\sigma|\tau_i$. Hence, $\sigma|\tau_i|\tau_j \in X^\#$, and thus condition 3. is met.

Next, we prove that $X^\#$ is the only evolutionarily stable set. Suppose that X is evolutionarily stable and that, contrary to the claim, there exists strategy profile $\sigma \in X$ such that $f_i(\sigma) < \alpha$ ($i = 1, 2$). We proceed in three steps to show that this leads to a contradiction. First, we construct strategy τ_i that behaves like σ_i against strategies in X , and is “nice” to the opponent. Second, we show that there exists τ_j ($j \neq i$) such that $\sigma|\tau_i|\tau_j \in X$. Third, we show that $\sigma|\tau_i|\tau_j$ is not a Nash equilibrium.

Step 1: For strategy $\sigma_j = (s_j, b_j)$, let τ_i be the associated modified strategy (t_i, b_i) , where the player i 's belief is the same as that of σ_i , where $t_i(b_i) = s_i(b_i)$, and where $t_i(a_j) = c$ for all $a_j \neq b_i$. In

other words, τ_i takes the same action with the same belief as σ_i does. Their payoffs are the same:

$$f_i(\sigma|\tau_i) = f_j(\sigma|\tau_i) = f_i(\sigma)$$

Step 2: When we define $\tau_j = \tau_i$, $f_j(\sigma|\tau_i|\tau_j) = f_j(\sigma|\tau_i)$ holds. By condition 3. in the definition of evolutionary stability, $\sigma|\tau_i|\tau_j \in X$.

Step 3: $\sigma|\tau_i|\tau_j$ is not a Nash equilibrium, since the strategy $\tau'_i = (t_i, c)$ by player i is a strictly better response to $\sigma|\tau_i|\tau_j$.

■

5 Concluding Comments

We have seen how the naive notion of transparency and reciprocal cooperation can be rescued, and how it is related to equilibrium selection.

The logic in our model is reminiscent of a “green beard.” The idea of a green-beard gene was proposed by Hamilton (1964) and named as “Green Beard” by Dawkins (1976). The concept remained a theoretical possibility until 1998, when a green beard gene was first found in nature, in the red fire ant (*Solenopsis invicta*). There is much literature in which “meta-players” are considered to be delegated people. One might consider that our model fits with a biological interpretation that “genes

delegate individuals.”

A setting such as infinitely repeated games justifies a number of individually rational outcomes, and tends to be out of place at equilibrium selection. However, by using the secret handshake argument in pre-play communication games, evolutionary stability selects only Pareto-efficient outcomes in our model.

Appendix

A.1 Example of Nonexistence of Nash Equilibrium

As will be appreciated from the folk theorem, our meta game often has many Nash equilibria. However the existence of a Nash equilibrium is not necessarily the case. For example, consider the following underlying game with no individually rational payoff profile:

		Player 2	
		c	d
Player 1	c	$1 + \varepsilon, -\varepsilon$	$-\varepsilon, 1 + \varepsilon$
	d	$\varepsilon, 1 - \varepsilon$	$1 - \varepsilon, \varepsilon$

where $\varepsilon > 0$ is sufficiently small. This underlying game is obtained by a perturbation of matching pennies. For any action profile, either player’s payoff is lower than her maximin payoff.

A.2 Mixed Extension

It is possible to extend the strategy sets in a meta game by permitting “mixed strategies,” symbolized by indices $*$. $s_i^* \in S_i^*$ is a probability distribution on S_i , and $b_i^* \in B_i^*$ is a probability distribution on B_i . We further define $\sigma_i^* = (s_i^*, b_i^*)$ etc.

Definition 5 Mixed strategy profile $\sigma^* = (s^*, \sigma^*)$ is feasible if for all s and b in the supports of s^* and b^* , respectively, (s, b) is feasible, and

$$\Pr \{b_i \text{ in } b_i^*\} = \sum_{(s,b) \in \Sigma: s_{-i}(b_{-i})=b_i} \Pr \{s_{-i} \text{ in } s_{-i}^*\} \Pr \{b_{-i} \text{ in } b_{-i}^*\},$$

for all i .

Lemma 2 In mixed Nash equilibrium $\sigma^* = (s^*, b^*)$,

$$f_i^*(\sigma^*) = f_i((s, b)),$$

for all s and b in the supports of s^* and b^* , respectively.

Example 6 We consider a 2×2 game:

		Player 2	
		c	d
Player 1	c	1, 1	0, 0
	d	0, 0	1, 1

There is a mixed Nash equilibrium in which both players play $cd | (pc + (1 - p) d)$ that means playing c if the opponent's action is c , playing d if the opponent's action is d , and believing that the opponent would take action c with probability p and d with $1 - p$. The outcome is (c, c) with probability p and (d, d) with $1 - p$.

A.3 More than Two-Person

We can easily extend the meta game to cases of three or more players. However the necessity part of the folk theorem does not hold. We present a counter example below.

Example 7 We consider a three-player game:

		Player 3					
		c			d		
		Player 2				Player 2	
		c	d			c	d
Player 1	c	3, 3, 0	1, 4, 0	Player 1	c	2, 2, 1	0, 3, 1
	d	4, 1, 0	2, 2, 0		d	3, 0, 1	1, 1, 1

The payoff structure for players 1 and 2 is similar to the above prisoners' dilemma game. In the underlying game, the unique dominant strategy equilibrium is (d, d, d) , and their maximin payoffs are 1. In the meta game, however, there is a Nash equilibrium (s, b) that induces the worst payoff for player 3, which is strictly less than her maximin payoff in the underlying game.

$$s_1(c, c) = s_1(c, d) = c,$$

$$s_1(d, c) = s_1(d, d) = d,$$

$$s_2(c, c) = s_2(d, d) = c,$$

$$s_2(c, d) = s_2(d, c) = d,$$

$$s_3(\cdot, \cdot) = c, \text{ and}$$

$$b_i = (c, c) \text{ for all } i,$$

where actions in parenthesis are arranged in ascending order of the players' names other than selves. As long as player 3 takes action c , the strategies of players 1 and 2 are the same as (Cd, Cd) in 2×2 games, and neither want to deviate from this situation. Although player 3 seems to have an incentive to take action d so as to obtain her maximin payoff 1 strictly better than present payoff 0, there is no feasible deviation for player 3 to take action d .

One may wonder why player 3 cannot play as she wishes. The author agrees that the collusion by players 1 and 2 not to let player 3 play d causes a feeling of strangeness. There exist other plausible equilibria, in one of which, for example, all play a meta extension of action d . Although we may feel ourselves compelled to exclude such Nash equilibria as impractical, doing so has a more serious effect. Consider the game below:

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

In this (underlying) game, player 1 is willing to play the same actions as player 2 plays. While player 2 is willing to play the same actions as player 1 plays if player 3 plays action c , she is willing to play different actions from player 1's actions if player 3 plays d . When we consider the meta game, the same strategy profile as above is a Nash equilibrium, which appropriately seems to reflect the incentives of players 1 and 2. Therefore, it would be overkill to exclude this strategy profile from the equilibrium concept. By comparison, there is no pure strategy Nash equilibrium in the underlying game.

References

- [1] Binmore, K., 1994. *Playing Fair: Game Theory and the Social Contract I*, The MIT Press.
- [2] Dawkins, R., 1976. *The Selfish Gene*, Oxford University Press.
- [3] Frank, R. H., 1987. "If Homo Economicus Could Choose His Own Utility Function, Would He Want One with a Conscience?" *American Economic Review*, **77**, 593–604.
- [4] Gauthier, D., 1986. *Morals By Agreement*, Oxford University Press.
- [5] Kim, Y.-G. and Sobel, J., 1995. "An Evolutionary Approach to Pre-Play Communication," *Econometrica* **63**, pp.1181-1193.
- [6] Kalai, A. T., E. Kalai, E. Lehrer, and D. Samet, 2010. "A Commitment Folk Theorem," *Games and Economic Behavior*, **69**, 127–137.

- [7] Keller, L, K. G. Ross, 1998. “Selfish Genes: a Green Beard in the Red Fire Ant,” *Nature* **394**, 573–575.
- [8] Levine, D. K., W Pesendorfer, 2007. “The evolution of cooperation through imitation,” *Games and Economic Behavior*, **58**, 293–315.
- [9] Hamilton, W. D., 1964. “The Genetical Evolution of Social Behaviour I–II,” *Journal of Theoretical Biology*, **7**, 1–52.
- [10] Howard, J. V., 1988. “Cooperation in the Prisoner’s Dilemma,” *Theory and Decision*, **24**, 203–213.
- [11] Howard, N., 1971. *Paradoxes of Rationality: Theory of Metagames and Political Behavior*. The MIT Press.
- [12] Robson, A. J., 1990. “Efficiency in Evolutionary Games: Darwin, Nash and Secret Handshake,” *Journal of Theoretical Biology* **144**, pp.379-396.
- [13] Tennenholtz, M., 2004. “Program Equilibrium,” *Games and Economic Behavior*, **49**, 363–373.
- [14] von Neumann, J., and O. Morgenstern, 1944. *Theory of Games and Economic Behavior*, Princeton University Press.
- [15] Wärneryd, K., 1991. “Evolutionary Stability in Unanimity Games with Cheap Talk,” *Economics Letters* **36**, 375–378.
- [16] Wärneryd, K., 2010. “Observable Strategies,” *mimeo*.