

Priors That Do Not Rule Out Strategic Uncertainty Cannot Lead to Nash Equilibrium¹

Eduardo Zambrano²

July 21, 2006

¹I would like to thank John Nachbar, Adam Brandenburger, Amanda Friedenberg, Francisco Rodríguez, Muhamet Yildiz, Oliver Board, Don Saari, Marek Kaminsky, David Levine, David Easley and Larry Blume for very useful conversations and participants at the 2006 Summer Meeting of the Econometric Society, the 2005 Spring Midwest Economic Theory Conference, at a UC-Irvine IMBS colloquium, and at two Notre Dame seminars for their comments to previous versions of this research. All remaining shortcomings are mine.

²Department of Finance, Mendoza College of Business, University of Notre Dame, Notre Dame, IN 46556. Phone: 574-631-4597, fax: 574-631-5255, e-mail: ezambran@nd.edu.

Abstract

Consider a two player game that is to be played once. The players receive information that they use to help them predict the choices made by each other. A *decision rule* for each player captures how each player uses the information received in making their choices. *Priors* in this context are probability distributions over the information that may be received and over the decision rules that their opponents may use. I investigate the existence of prior beliefs for each player that satisfy the following properties: (R) they do not rule out their opponent using a rational decision rule, (K) they do not rule out the existence of information that that would reveal the choice made by the opponent, and (SU) they do not rule out *strategic uncertainty*—a belief diversity condition. In this paper I show that for a large class of games there are no prior beliefs that satisfy properties (R), (K) and (SU). In the paper I discuss the implications of this result, in particular regarding whether one should expect a Nash equilibrium to arise in a game that is to be played once.

1 Introduction

Imagine that two individuals enter a room and are given information about a game they are about to play with each other, once. The purpose of this paper is to provide an answer to the following question: can we expect these players to play a Nash equilibrium of the game?

It is difficult to overstate the importance of this question as Nash equilibrium is the central solution concept used at making predictions about what will transpire in a non-cooperative game that is played once. It may suffice here to recall Myerson's (1999) initial words in his article on Nash equilibrium and the history of economic theory: "The formulation of Nash equilibrium has had a fundamental and pervasive impact in economics and the social sciences which is comparable to that of the discovery of the DNA double helix in the biological sciences." (p. 1067)

Many applications of game theory assume players will play a Nash equilibrium of a non-cooperative game even if the game is only going to be played once. However, there is substantial experimental evidence that in such situations play deviates systematically from equilibrium. This has been prominently exemplified in the work by Nagel (1995), Ho, Camerer, and Weigelt (1998), Stahl and Wilson (1994, 1995), Costa-Gomes, Crawford and Broseta (1998, 2001), Costa-Gomes and Weizsäcker (2005), Camerer, Ho and Chong (2004) and Costa-Gomes and Crawford (2006).

In general, the experimental literature appeals to some form of bounded rationality to explain these findings. In this paper I put bounded rationality considerations aside and explore whether there is a conceptual barrier as to why play systematically deviates from equilibrium in one shot games. In other words, whereas the experimental evidence shows that conditions that lead to Nash equilibrium are difficult to satisfy in practice, I ask in the present paper whether the conditions that lead to equilibrium can even obtain *in principle*. An affirmative answer to this question would allow us to see the experimental evidence in a new light, as it would show that the results that have been obtained in the experiments are *natural and to be expected*, given the intrinsic difficulties players have in reasoning their way to a Nash equilibrium, rather than *pathological and susceptible to go away* if one only found the right experimental treatment in which to examine Nash's theory.

From a conceptual standpoint, the state of the art regarding conditions that lead to Nash equilibrium is developed in the seminal paper by Aumann and Brandenburger (1995), which identified conditions about what the play-

ers must know and believe about each other for the actions chosen by the players or the beliefs held by them to be a Nash equilibrium. In their *Preliminary Observation* they notice that rationality plus mutual knowledge of the action profile chosen by all players suffices for that action profile to be a Nash equilibrium. In addition, in their *Theorem A* (for two player games) they show that mutual knowledge of rationality, mutual knowledge of the structure of the game, and mutual knowledge of the conjectures the players have about each other's choices suffices for those conjectures to be a Nash equilibrium. Also, in their *Theorem B* (for n player games) they show that if the players have a common prior, their payoff functions and their rationality are mutually known, and that their conjectures are commonly known, then for each player j , all the other players agree on the same conjecture about j , and the resulting profile of conjectures is a Nash equilibrium.¹

In this paper I show that there is indeed a conceptual barrier to interpreting the Aumann and Brandenburger (1995) results as providing support for the notion that one should expect Nash equilibrium to arise in a one shot play. This barrier stems from the fact that the conditions that lead to Nash equilibrium in Aumann and Brandenburger (1995) are difficult to obtain if they are not assumed at the outset. In this paper I develop a very simple formalism to show that this is indeed the case.

An informal description of the setup in which I operate is as follows. Imagine that two players are about to enter a room, or *context*, in which they simultaneously receive some information m regarding a game that they are about to play with each other. After this, each player independently chooses an action from their action set.

A *decision rule* for any such player is a map from the information revealed to the players to an action set. Therefore, before the players enter the room they have beliefs about the decision rules chosen by their opponents and about the information m to be received. One could assume, as in a Bayesian Nash equilibrium, that these beliefs are *degenerate*, that is, that they are point mass on the actual decision rule chosen by the opponent. Naturally, a degenerate prior rules out strategic uncertainty about how the opponent will play conditional on the information that is received. I stay away from this assumption, as this would beg the question as to how did

¹For excellent surveys on these and other results on the epistemic foundations of game theory see Bonanno and Battigalli (1999), Brandenburger (1999, 2006), Dekel and Gul (1997) and Samuelson (2004).

these beliefs configure themselves in equilibrium proportions to begin with. In this paper I work with prior beliefs that are nondegenerate and hence not concentrated on the actual decision rules chosen by each opponent. I label these as *priors that do not rule out strategic uncertainty*.

I am therefore interested in investigating the existence of priors that do not rule out strategic uncertainty that, conditional on the information that is revealed at the context, would lead the players' knowledge and belief to satisfy the Aumann and Brandenburger conditions for Nash equilibrium.

More specifically, I wish to impose the following properties on these prior beliefs:

(R). That they are *consistent with rationality*: Player 1 does not rule out that 2 uses a rational decision rule. Similarly for player 2.

(K). That they *can lead to mutual knowledge of actions*: Player 1 does not rule out that he may receive information that would reveal the choice that player 2 makes. Similarly for player 2.

(SU) That they *do not rule out strategic uncertainty*: if player 1 does not rule out a decision rule for player 2 then there is a simple variant of this decision rule that is a decision rule for player 1, and that player 2 does not rule out for player 1.

Loosely, the main result of the paper, Theorem 1, is that for a large class of games there are no prior beliefs for players 1 and 2 that satisfy R, K and SU.²

The interpretation of this result is as follows: mutual knowledge of actions cannot arise as a consequence of the players having access to some common information about the game to be played in a setup where players's prior beliefs do not rule out strategic uncertainty, namely, when the players are both equally cautious about how each other respond to the information received. This poses a barrier to the interpretation of Aumann and Brandenburger's *Preliminary Observation* as providing support for the notion that players would play a Nash equilibrium of a non-cooperative game that is to be played just once.

It is quite fair to say that the real contributions in Aumann and Brandenburger (1995) are their Theorems A and B, which make no assumptions about whether the action profile chosen by the players is mutual knowledge.

²Theorem 1 is actually stronger than this since it works with ε -rationality and mutual p -belief of action profiles. I have omitted those details in this introduction to aid exposition.

Unfortunately it turns out that mutual knowledge of action profiles is *implied* by the assumptions in Theorems A and B for a large class of games, namely, those with strict Nash equilibria. This is Theorem 5 in the present paper. This means that for a large class of games the assumptions in Theorems A and B cannot be satisfied if mutual knowledge of actions does not hold. In light of Theorem 1, this poses a barrier to the interpretation of Aumann and Brandenburger's *Theorem's A and B* as providing support for the notion that the players' conjectures would be a Nash equilibrium of a non-cooperative game that is to be played just once.

The structure of the rest of the paper is as follows: In Section 2 I develop an example that illustrates the main issues and results. In Section 3 I develop the model of a non-cooperative game in which I prove my results. In Section 4 I state and prove the main results. Finally, in Section 5, I discuss further the interpretation of the results together with the related literature.

2 An Illustration

Consider the two player game G :

		<i>Bob</i>	
		<i>c</i>	<i>d</i>
<i>Ann</i>	<i>C</i>	3, 3	0, 0
	<i>D</i>	0, 0	1, 1

and a situation in which both players receive one of two pieces of information, m_1 or m_2 , before the game is played. Each player has beliefs about the information that they will receive, and about what this information may say about the action chosen by the opponent.

(i) Consider a case in which each player is certain that the opponent uses *constant* decision rules, namely, rules that choose the same action regardless of the information received. An example of what these beliefs would look like is described in CONTEXT 1 below:

$\boxed{\beta_A:}$

	<i>cc</i>	<i>cd</i>	<i>dd</i>	<i>dc</i>
m_1	49%	0	1%	0
m_2	1%	0	49%	0

$\boxed{\beta_B:}$

	<i>CC</i>	<i>CD</i>	<i>DD</i>	<i>DC</i>
m_1	49%	0	1%	0
m_2	1%	0	49%	0

Read the example as follows: β_A is Ann's belief over the decision rules of Bob and over the information the players may receive. A decision rule for Bob captures how Bob chooses an action from his action set based on the information received. For example, decision rule cd for Bob is one in which he chooses c if he receives information m_1 and chooses d if he receives information m_2 . The numbers inside the cells denote the *ex ante* probabilities that the players ascribe to the information and the decision rule corresponding to that cell. For example, according to β_A , Ann puts probability 1% on receiving information m_1 and on Bob using decision rule dd . Notice also that Ann puts probability zero on Bob using rules cd and dc . Similar statements can be made about Bob's prior belief β_B .

In this example Ann considers possible only two decision rules for Bob: one according to which Bob plays c regardless of the information received (decision rule cc), and one according to which Bob plays d regardless of the information received (decision rule dd). Ann's beliefs about Bob's choices are also determined by β_A : when the information received by the players is m_1 Ann puts probability of 98% on Bob choosing c , whereas when the information received by the players is m_2 Ann puts probability of 2% on Bob choosing c . Similar statements can be made regarding Bob's beliefs.

Notice that for every decision rule Ann considers possible for Bob there is information that reveals with high probability the choice made by that decision rule of Bob: for rule cc the information is m_1 , and for rule dd the information is m_2 .³ Similarly for every decision rule Bob considers possible for Ann. These beliefs therefore satisfy property K. Notice also that the beliefs for both players are rather similar regarding the kinds of decision rules they consider possible for each other. These beliefs satisfy property SU.⁴

The starting point of this paper is the observation that the support of each player's beliefs are inconsistent with rationality (therefore violating property R). This is so because each player is certain that her opponent uses a decision rule that is independent of the information received even though, if

³Ann's belief over c conditional on m_1 is 98% and Bob in fact chooses action c when using decision rule cc and receives information m_1 . In turn, Ann's belief over d conditional on m_2 is 98% and Bob in fact chooses action d when using decision rule dd and receives information m_2 .

⁴This is, of course, an informal statement. The example exactly satisfies properties SU and K, as these properties are defined in Section 3 of the paper.

players actually use rational decision rules, their decision rules will not be independent of the information received: given these beliefs, a rational decision rule for Ann would choose C if the information received is m_1 and D if the information received is m_2 (decision rule CD); similarly for Bob and rule cd .

The situation is complicated if both players use rational decision rules and somehow the structure of the game and the belief structure is known to both players for then Ann, correctly knowing Bob's payoffs, and also knowing that Bob puts positive probability only on the constant rules for Ann, would be certain, *and wrong*, that Bob does not use a rational decision rule in making his choice. Similarly for Bob.

(ii) One could try to remedy this situation by adding to the supports of the player's beliefs the decision rules that vary with the information received. Then, for this game, beliefs would not rule out strategic uncertainty (would satisfy **SU**) and would be consistent with rationality (they would satisfy **R**) but then *they could not lead to mutual knowledge of actions* (they would violate **K**). To see this assume that the beliefs are now given by **CONTEXT 2** below:

β_A :

	cc	cd	dd	dc
m_1	24%	24%	1%	1%
m_2	1%	24%	24%	1%

β_B :

	CC	CD	DD	DC
m_1	24%	24%	1%	1%
m_2	1%	24%	24%	1%

Now Ann's beliefs, conditional on m_1 are that with probability 96% Bob chooses c . Conditional on m_2 , Ann believes that Bob chooses c with probability 4%.

Clearly, the beliefs do not rule out strategic uncertainty and include a rational decision rule for each player in their supports (hence satisfying **SU** and **R**).

Now let's check whether **K** holds, that is, whether for every decision rule Ann considers possible for Bob there is information that reveals with high probability the choice made by that decision rule of Bob: for rule cc the information is m_1 , for rule dd the information is m_2 , and for rule cd either piece of information, m_1 or m_2 , works. Now notice that for rule dc there is no information that would reveal with high probability the choice made by that decision rule of Bob, since if the information is m_1 rule dc picks d

and Ann believes with probability 96% that Bob chooses c . In turn, if the information is m_2 rule dc picks c and Ann believes with probability 96% that Bob chooses d .

(iii) One could, again, try to remedy this situation by adding more pieces of information to the information space (e.g, information that would identify dc). But as one includes more information the set of decision rules for each player also changes, therefore making the rational decision rule more complicated. For example, imagine now that there are now three pieces of information, m_1 , m_2 and m_3 , (where m_3 is a ‘copy’ of m_2 in that players are certain that each other chooses at m_3 as they did at m_2) and that the beliefs are as in CONTEXT 3 below

$\beta_A:$					$\beta_B:$				
	ccc	cdd	ddd	dcc		CCC	CDD	DDD	DCC
m_1	23%	23%	1%	1%	m_1	23%	23%	1%	1%
m_2	1%	23%	23%	1%	m_2	1%	23%	23%	1%
m_3	2%	0	0	2%	m_3	2%	0	0	2%

Again, in this case, it is not hard to see that the beliefs have supports that are structurally similar, therefore satisfying SU. It is also the case that for every decision rule Ann considers possible for Bob there is information that reveals, with high probability, the choice made by Bob under that rule, therefore satisfying K. To see this notice that for rule ccc the information is m_1 , for rule ddd the information is m_2 , for rule cdd either piece of information, m_1 or m_2 , works, and for rule dcc , the information is m_3 . Yet, again, in this example, R fails: given these beliefs, Bob’s rational decision rule is cdc , which Ann has ruled out for Bob in this case. Then, if players are rational, beliefs over decision rules will be *mis-specified*.

(iv) One could, once more, try to remedy this situation by altering the beliefs in some other way. *The main message of this paper is that there is no remedy to this situation:* there are no beliefs for this game for which R, K and SU simultaneously hold.

Theorem 1 in the present paper shows that the inconsistency between R, K and SU that these examples display is not specific to the particular game being considered in this Section but that it indeed holds for a very large class of games, namely, for those games without a weakly dominant action for each player.

3 The setup

3.1 The context of a non-cooperative game

I consider standard finite two player⁵ games $G = \langle A_1, A_2, g_1, g_2 \rangle$. I define $\mathcal{C} = \langle M, \beta_1, \beta_2 \rangle$, a *context* for game G , as a finite information space M and a prior belief β_i for each player i . The interpretation is that first the players receive some common information (information m , from the information space M) and then simultaneously make their action choices.

Throughout the paper I remain agnostic as to the content of this information. It is not payoff relevant but it may, or may not, contain information about the opponent's payoffs, actions, or beliefs.

Given M one can define a *decision rule* for player i as a map $r_i : M \rightarrow A_i$. Any such decision rule encodes how the information revealed will affect the choices made by each player. Let R_i be the set of decision rules for player i .

Players do not know which decision rule their opponents are employing for the purpose of making their choice of action. Instead they have beliefs over which rule it may be. Hence, player i 's prior belief β_i is a probability distribution over the set $R_{-i} \times M$. Let \hat{R}_{-i} be the support of player i 's belief β_i over R_{-i} , let $\hat{R} = \hat{R}_1 \times \hat{R}_2$ and let $\beta_i^a(\cdot; m)$ be the belief over A_{-i} held by player i that is induced by the conditional distribution $\beta_i(\cdot | m)$. Therefore, fixing \mathcal{C} not only determines M and β_i but also \hat{R} and β_i^a .

I am interested in whether a context \mathcal{C} exists that satisfies certain natural properties. I want the context to be consistent with rationality, to permit mutual knowledge of actions and not to rule out strategic uncertainty. Formal definitions of these properties follow.

3.2 The rationality of a context

Fix $\varepsilon \geq 0$. Given β_1 and g_1 , rule r_1 in R_1 is an ε -rational decision rule for player 1 if, for every $m \in M$:

$$\sum_{a_2 \in A_2} \beta_1^a(a_2; m) g_1(r_1(m), a_2) + \varepsilon \geq \sum_{a_2 \in A_2} \beta_1^a(a_2; m) g_1(a_1, a_2) \text{ for all } a_1 \in A_1$$

⁵I focus on two player games, but I see no difficulty in extending the result to games with more than two players.

Similarly for player 2.

Context \mathcal{C} is *consistent with ε -rationality* if for each player i there is at least one decision rule r_i in \widehat{R}_i that is an ε -rational decision rule for player i . If this holds I say that the context satisfies property R. The interpretation is that each player does not rule out the opponent using a rational decision rule. It is a very weak assumption.

3.3 Contexts that can lead to mutual “knowledge” of actions

Fix $p \in (0, 1]$. Context \mathcal{C} *can lead to mutual p -belief of actions* if for every $r_2 \in \widehat{R}_2$ there is some information $m \in M$ such that

$$\beta_1^a(r_2(m); m) \geq p,$$

and similarly with the roles of the players reversed. If this holds I say that the context satisfies property K.

The interpretation is as follows. Assume for a moment that player 1 picks his decision rule from among those that player 2 has not ruled out for him. Property K requires that there is a piece of information m that, were player 2 to receive it, it would reveal (with high probability) the actual action chosen by player 1 according to this information and his chosen rule.

To actually see how weak condition K is assume that a pair of players are about to play game G in context \mathcal{C} . Assume that for the purpose of deriving their choice in G player 1 uses rule r_1 and player 2 uses rule r_2 . Assume that the players receive information $m \in M$. Even if this context satisfies property K this has no implications whatever on the beliefs each player has, conditional on m , over each other's choices. All it says is that there must have been some information m' and m'' , not necessarily the same for each player, that the players could have received that would have led the players to assign high probability to the actual actions chosen by each other when their opponents also received such information m' and m'' .

It is also instructive to consider a stronger version of K, dubbed K', where, for each strategy profile in the support of the player's beliefs, the information that reveals the player's choices is the same across players. Fix $p \in (0, 1]$. Context \mathcal{C} satisfies property K' if for any $(r_1, r_2) \in \widehat{R}$ there is information

$m \in M$ such that

$$\begin{aligned}\beta_1^a(r_2(m); m) &\geq p \text{ and} \\ \beta_2^a(r_1(m); m) &\geq p.\end{aligned}$$

The main result in the present paper uses the weaker property K. Property K' plays a role in Section 5, where I compare my results with those of Aumann and Brandenburger (1995).

3.4 Strategic Uncertainty

Let $\gamma_{12} : A_1 \rightarrow A_2$ be a function that ‘translates’ actions from the action space of player 1 to the action space of player 2.

Rule r_2 for player 2 is a *mechanical brother of rule r_1* if there is a translation γ_{12} such that

$$r_2(m) = \gamma_{12}(r_1(m)) \text{ for every } m \in M.$$

Similarly for player 1.

Context \mathcal{C} *does not rule out strategic uncertainty* if every mechanical brother r_2 of $r_1 \in \widehat{R}_1$ is in \widehat{R}_2 ; similarly for player 1. If this holds I say that the context satisfies property **SU**.

The intention behind requiring that beliefs do not rule out strategic uncertainty is to ensure that beliefs have support that are sufficiently diverse (in particular, that they are not concentrated around the actual decision rules being employed by the players). I make this assumption because it makes little sense to investigate how is it that actions can become mutual knowledge in a setup in which the decision rules are mutual knowledge themselves. One would then have to explain how this came to be the case. In other words, I do not wish to impose equilibrium-like restrictions on the beliefs over decision rules as I try to explain how equilibrium in actions (or in beliefs over actions) arises as a result of deliberation and conditioning on the information that may be received about a game to be played in a static setting.

Remark 1 *It is important to keep in mind that beliefs do not have to have full support on the entire set of decision rules of the opponent to satisfy **SU**. For example, the beliefs in Contexts 1 and 3 in Section 2 are not full support on the set of decision rules, yet they satisfy **SU**.*

The discussion above reveals that assuming **R**, **K** or **SU** in isolation is not very restrictive. The situation is very different if one wishes to assume them jointly.

4 The results

An action a_1^* is weakly dominant if, for any $a_2 \in A_2$,

$$g_1(a_1^*, a_2) \geq \max_{a_1 \in A_1} g_1(a_1, a_2).$$

The definition for player 2 is similar.

Game G satisfies No Weak Dominance (NWD) if neither player has a weakly dominant action.

The main result of the paper is the following:

Theorem 1 *Suppose that NWD holds in game G . Then there is $\bar{\varepsilon} > 0$ and $\bar{p} \in (0, 1]$ such that, for any $\varepsilon \in [0, \bar{\varepsilon})$ and any $p \in (\bar{p}, 1]$ there is no context \mathcal{C} for G that satisfies **R**, **K** and **SU**.*

The interpretation is that, for a large set of games, if beliefs do not rule out strategic uncertainty and can lead to mutual p -belief of actions then they are inconsistent with the opponent employing an ε -rational decision rule. Then, to the extent that players employ a rational decision rule when facing game G and context \mathcal{C} , then the players beliefs *over decision rules* will necessarily be mis-specified.

Remark 2 *The result breaks down for games with a weakly dominant action for each player because beliefs that do not rule out strategic uncertainty would always be consistent with rationality as \hat{R} would contain the decision rule for each player that always chooses the weakly dominant action for that player.*

The proof of Theorem 1 relies on the notion of an evil twin, which I define below.

4.1 Evil twins

As in Nachbar (1997, 2005) define the function $\tilde{a}_2 : A_1 \rightarrow A_2$ for any action $a_1 \in A_1$ by

$$\tilde{a}_2(a_1) \in \arg \max_{a_2 \in A_2} \left[\max_{a'_1 \in A_1} g_1(a'_1, a_2) - g_1(a_1, a_2) \right].$$

Loosely, when player 2 chooses $\tilde{a}_2(a_1)$ player 1 has maximal incentives not to choose a_1 .

Rule $r_2 \in R_2$ is an *evil twin* of $r_1 \in R_1$ if

$$r_2(m) = \tilde{a}_2(r_1(m)) \text{ for every } m \in M$$

A similar definition holds for player 2.

Context \mathcal{C} has the *evil twin* property if for any decision rule $r_1 \in \hat{R}_1$ there is a decision rule $r_2 \in \hat{R}_2$ that is an evil twin of r_1 .

A decision rule r_1 cannot be rational for player 1 if there is some information m that would lead player 1 to know the choice made by player 2 and player 2's decision rule is an evil twin of r_1 . The following lemma, adapted from a result by Nachbar (1997) on Bayesian learning, records this fact. The proof is in the Appendix.

Lemma 2 *Suppose that NWD holds in game G . Then there is $\bar{\varepsilon} > 0$ and $\bar{p} \in (0, 1]$ such that, for any $\varepsilon \in [0, \bar{\varepsilon})$ and $p \in (\bar{p}, 1]$, and any context \mathcal{C} of G , and any $r_1 \in R_1$, if \tilde{r}_2 is an evil twin of r_1 and $\beta_1^a(\tilde{r}_2(m); m) \geq p$ for some $m \in M$ then r_1 is not an ε -rational decision rule for player 1. Similarly for player 2.*

The next step in the argument is to observe that evil twins of the sort constructed above will be included in the support of priors that satisfy SU.

Lemma 3 *Let G be a finite two player game and \mathcal{C} be a context for G . If \mathcal{C} satisfies SU then it also satisfies the evil twin property.*

Proof. Pick any $r_1 \in \hat{R}_1$. Define $\gamma_{12} : A_1 \rightarrow A_2$ by $\gamma_{12}(a_1) = \tilde{a}_2(a_1)$. It then follows from condition SU that \hat{R}_2 contains an evil twin of r_1 and that \mathcal{C} satisfies the evil twin property. ■

I am now in a position to prove the paper's main result.

4.2 The proof of Theorem 1

Fix $\bar{\varepsilon}$ and \bar{p} as in Lemma 2. Pick $\varepsilon < \bar{\varepsilon}$ and $p \in (\bar{p}, 1]$. Assume there is a context \mathcal{C} for G that satisfies K and SU. I will show that \mathcal{C} violates R.

Pick any $r_1 \in \hat{R}_1$. Since \mathcal{C} satisfies SU, Lemma 3 implies that there is $\tilde{r}_2 \in \hat{R}_2$ that is an evil twin of r_1 . By property K, there is $m \in M$ such that

$\beta_1^a(\tilde{r}_2(m); m) \geq p$. Then, by Lemma 2, r_1 is not an ε -rational decision rule for player 1. Since r_1 was arbitrary it follows that \mathcal{C} is not consistent with ε -rationality. A similar argument can be built starting with any $r_2 \in \hat{R}_2$. ■

For the interpretation of this result see the Introduction and Section 5.

4.3 Games with strict Nash equilibria

A *strict* Nash equilibrium is a Nash equilibrium in which each player would do worse if she were to deviate from the action prescribed to her by the Nash equilibrium profile. Strict Nash equilibria arise, for example, when the best response correspondences of the players are single valued. Many games of interest have only strict Nash equilibria.

Let $G^s = \langle N, (A_i), (g_i^s) \rangle$ be an n -player game with only strict Nash equilibria. Let \mathcal{G}_i be a set of payoff functions for player i and consider an epistemic model such as the one developed in Aumann and Brandenburger (2005). Such *epistemic model* is essentially the profile $\langle (T_i)_{i \in N} \rangle$, where T_i is a set, each element of T_i is a probability distribution over $A_{-i} \times \mathcal{G}_{-i} \times T_{-i}$ and the states of the world are members of $S = \times_{i \in N} A_i \times \mathcal{G}_i \times T_i$. The interpretation is that each state of the world specifies, for each player, the action chosen by the player, the player's payoff function, and the player's beliefs over the opponents' action, payoff and belief.

It turns out that the following is true:

Proposition 4 *Let σ be an n -tuple of mixed strategies. Suppose that at some state of the world s it is mutually known that the players are all rational, that the game G^s is being played, that the conjecture of each player i about each other player j is σ_j , and that it is independent of i 's conjecture about all other players. Then the action profile chosen by the players is mutually known at s .*

Proof. By Lemma 2.6 in Aumann and Brandenburger (1995) it follows that the players are indeed rational, that the game being played is actually G^s and that the conjecture of each player i about each other player j is in fact σ_j . By Remark 7.1 in Aumann and Brandenburger (1995) σ is a Nash equilibrium of G^s . By the definition of G^s , σ is a strict Nash equilibrium, hence, for each player j , σ_j puts probability one on a single action a_j . Each player i then puts probability one on player j being rational, on j 's payoff being g_j^s , and on j 's conjecture about all other players being (a_{-j}) . Then player i puts

probability one on j choosing the only action that maximizes g_j^s with respect to (a_{-j}) . By the definition of G^s , this action is equal to a_j , which is indeed the action that player j chooses at s . ■

The following result is a consequence of combining the content of Theorems A and B in Aumann and Brandenburger (1995) with Proposition 4.

Theorem 5 *For games with only strict Nash equilibria the hypotheses in Theorem A and Theorem B in Aumann and Brandenburger (1995) about what must be true at a certain state of the world s imply mutual knowledge of the action profile chosen by the players at s .*

Proof. First consider the hypotheses in Theorem A (two players) for game G^s . These are: mutual knowledge of rationality, of conjectures, and of G^s . Since there are only two players, independence is automatically satisfied and Proposition 4 implies that the action profile chosen at s is mutual knowledge. Next consider the hypotheses in Theorem B for game G^s . These are: a common prior, mutual knowledge of the game being G^s and of the player's rationality, and that their conjectures are commonly known. Then Theorem B implies that, for each player j , all the other players agree on the same conjecture about j and the proof of Theorem B implies that these conjectures are independent. Then Proposition 4 implies that the action profile chosen at s is mutual knowledge. ■

This implies that for games with only strict Nash equilibria the assumptions in Theorems A and B cannot be satisfied in a context where mutual knowledge of actions cannot be satisfied. More on this in the Section below.

5 Discussion

5.1 Interpretation of the results

Let G be a game and let p and ε be as in Theorem 1. In addition, let \mathcal{C} be a context for G that satisfies K and SU. Then, by Theorem 1, the supports of the player's beliefs β_i must not contain an ε -rational decision rule for either player. The implication is that if both players employ rational decision rules when facing game G and context \mathcal{C} then beliefs over decision rules at \mathcal{C} will be misspecified as each player would be using a rule that necessarily would be outside of the support of the opponent's belief. Therefore, under the

hypotheses of Theorem 1, the players can use rational decision rules or have well-specified beliefs over decision rules, but not both. *Under belief diversity, if players use rational decision rules and have beliefs consistent with mutual knowledge of actions, then the players will by necessity be wrong about each other's decision rules.*

Actually, the situation is worse than this suggests. Let G be a game and let p and ε be as in Theorem 1. In addition, let \mathcal{C} be a context for G that satisfies **K** and **SU**. Assume that both players use ε -rational decision rules. Suppose that player 1 somehow knows correctly that \mathcal{C} satisfies **K** and **SU** and that also correctly knows player 2's payoff function. *It follows that player 1 is certain (and wrong!) that player 2 is not using an ε -rational decision rule.* This is so because player 1, correctly knowing G and that \mathcal{C} satisfies **K** and **SU**, can use Theorem 1 to conclude that \hat{R}_2 does not contain any ε -rational decision rule for player 2. By the definition of \hat{R}_2 , this means that player 1 is certain that player 2 is not using an ε -rational decision rule. Similarly for player 2.

One is then left with a formalism populated by rational decision makers that are certain that their opponents aren't rational decision makers themselves! Not that one would want this, but this is precisely the point: it is what would follow from the seemingly innocuous assumptions of rational decision, mutual knowledge of payoffs, and mutual knowledge of context \mathcal{C} satisfying **K** and **SU**.

5.2 Relation to Aumann and Brandenburger (1995)

The results presented here do not contradict in any way those of Aumann and Brandenburger (1995). Needless to say, those results and their proofs are correct as stated. Moreover, the assumptions behind the Preliminary Observation, and Theorems A and B are not vacuous, as the discussion in their Section 7.h shows. The point that I am trying to make in the present paper is different: that those assumptions are difficult to obtain if they are not assumed at the outset.

Perhaps the best way to understand this, given a two player game G and information space M is to consider an epistemic model for G when G is viewed as a game "with a move by Nature," namely, to consider the epistemic model $\langle T_1^M, T_2^M \rangle$ where every element $t_i^M \in T_i^M$ is a probability distribution over $R_{-i} \times M \times \mathcal{G}_{-i} \times T_{-i}^M$. In this setting, I define an *ex-ante state of the world* as a member of $W = R_1 \times \mathcal{G}_1 \times T_1^M \times R_2 \times \mathcal{G}_2 \times T_2^M$ and an *interim state of*

the world as a member of $S = W \times M$.

The interpretation is that each *ex ante* state of the world specifies the payoff functions for the players, what decision rule each player will use for choosing an action after information m is revealed, and the player's beliefs over M and over each other's rules, payoffs and beliefs. An *interim* state of the world, in addition, reveals such information m to the players and therefore it also specifies the profile of actions chosen by the players at that state. An *interim state* is essentially identical to a *state* in the conventional epistemic model of Aumann and Brandenburger (1995).

Notice also that, given $\langle T_1^M, T_2^M \rangle$, any *ex ante* state $\omega \in W$ induces a context $\mathcal{C}[\omega] = \langle M, \beta_1^{[\omega]}, \beta_2^{[\omega]} \rangle$ for the game $G[\omega] = \langle A_1, A_2, g_1^{[\omega]}, g_2^{[\omega]} \rangle$, where $\beta_i^{[\omega]}$ is defined to be player i 's belief over $R_{-i} \times M$ at ω and $g_i^{[\omega]}$ is the payoff function for player i at ω . I say that *ex ante* state $\omega \in W$ satisfies **R**, **SU** or **K** whenever $\mathcal{C}[\omega]$, respectively, satisfies **R**, **SU** or **K** for $G[\omega]$.

Let $r_i^{[\omega]}$ be the decision rule that player i uses at $\omega \in W$ and let $\hat{R}_i[\omega]$ be the support of player i 's beliefs over at ω . *Rationality* at ω for player i can then be defined in a similar way, namely, when such $r_i^{[\omega]}(m)$ is a best response to $\beta_i^{a[\omega]}(;m)$ for all $m \in M$. Priors are *well-specified* at ω if $r_i^{[\omega]} \in \hat{R}_i[\omega]$ for all players. *Rationality* at an interim state $s = (\omega, m)$ for player i simply refers to the action $r_i^{[\omega]}(m)$ being a best response to $\beta_i^{a[\omega]}(;m)$ for the information m received by the players at s .

I say that *priors can lead to Nash equilibrium* at ω if priors at ω are well-specified and satisfy **K'** (for $p = 1$).

This definition makes sense because if the players are rational at ω and the priors can lead to Nash equilibrium at ω this means that there is an interim state $s = (\omega, m)$ where both players are rational and have mutual knowledge of action profiles,⁶ hence making the actions chosen at s a Nash equilibrium for the game being played. Moreover, if the game being played at ω has only strict Nash equilibria and mutual knowledge of actions was not possible at any state $s = (\omega, m)$ then it follows from Theorem 5 that the epistemic conditions for beliefs to be a Nash equilibrium from either Theorem A or B in Aumann and Brandenburger (1995) could not be satisfied at any such s .

⁶To see this notice that, since priors at ω are well-specified then, by property **K'**, there is $m \in M$ such that $\beta_i^a(r_{-i}^{[\omega]}(m); m) = 1$ for every player i .

With this in place one can see what bearing Theorem 1 has on whether the epistemic conditions that lead to Nash equilibrium can arise at some interim state $s = (\omega, m)$ that has priors that do not rule out strategic uncertainty at ω .

Theorem 6 *Suppose that NWD holds in game G . Let $\omega \in W$ be an ex-ante state of the world of model $\langle T_1^M, T_2^M \rangle$ where G is being played. If players are rational and their priors do not rule out strategic uncertainty at ω then those priors cannot lead to Nash equilibrium at ω .*

Proof. Assume that there is an *ex ante* state ω such that priors *can* lead to Nash equilibrium at ω , that is, assume that priors at ω are well-specified and satisfy K' (and therefore K) for $p = 1$. Since the priors do not rule out strategic uncertainty at ω then by definition they satisfy SU . Then, by Theorem 1 the priors at ω do not satisfy R (for $\varepsilon = 0$), that is, they rule out the players being rational. Since, by assumption, the players *are* rational at ω then the priors must be mis-specified at ω . The contradiction establishes the result. ■

It is important to stress that Theorem 6 does not say that the Aumann and Brandenburger conditions cannot be satisfied in any interim state of model $\langle T_1^M, T_2^M \rangle$. They can, but under CS and K' this will happen as the players' beliefs are mis-specified about the decision rules of the opponent, if the players use rational decision rules to begin with. The example below shows how this is so.

Consider again the game G discussed in Section 2:

		<i>Bob</i>	
		<i>c</i>	<i>d</i>
<i>Ann</i>	<i>C</i>	3, 3	0, 0
	<i>D</i>	0, 0	1, 1

and the following context for G (CONTEXT 4):

$\beta_A:$	<i>cccc</i>	<i>dddd</i>	$\beta_B:$	<i>CCCC</i>	<i>DDDD</i>
m_1	25%	0	m_1	25%	0
m_2	0	25%	m_2	0	25%
m_3	25%	0	m_3	0	25%
m_4	0	25%	m_4	25%	0

This context is similar to CONTEXT 1 in Section 2 in that the players put positive probability only on the constant decision rules for the opponent. It is not hard to see that this context satisfies **SU**, as each player considers possible *all* of the constant rules for the opponent. CONTEXT 4 also satisfies **K'** since for the profile $(CCCC, cccc)$ information m_1 reveals the choices made by the players when they receive information m_1 , and similarly for the profiles $(DDDD, dddd)$, $(DDDD, cccc)$, and $(CCCC, dddd)$ when the information received by the players is, respectively, m_2 , m_3 , or m_4 . Now, given these beliefs, Ann's rational decision rule is $CDCD$ and Bob's rational decision rule is $cddc$. Because these rules are outside of the support of the beliefs of the players, CONTEXT 4 violates **R**.

Now with these ingredients build an epistemic model $\langle T_1^M, T_2^M \rangle$ for G adapted to CONTEXT 4, namely, let $T_A^M = \{\beta_A\}$, $T_B^M = \{\beta_B\}$, where β_A and β_B are as in CONTEXT 4 above. The interpretation is that in this example there is only uncertainty about R and M ; not about T^M or \mathcal{G} . Consider the *interim* state $s = (CDDC, g_A, \beta_A, cdcd, g_B, \beta_B, m_1)$. It is not hard to see that both players are rational at s , and that the action profile is mutual knowledge at s . By Aumann and Brandenburger's Preliminary Observation, the action profile chosen at s must be a Nash equilibrium. Such profile is (C, c) , indeed a Nash equilibrium of G .

It turns out that at this state s rationality (of actions), payoffs and conjectures are mutual knowledge. Hence, by Theorem A in Aumann and Brandenburger, the conjectures the players have about each other must be a Nash equilibrium. At state s Ann's conjecture about Bob is that with probability one Bob plays c ; Bob's conjecture about Ann is that with probability one Ann plays C . Those conjectures are indeed a Nash equilibrium of G . Moreover, G has only strict Nash equilibria, and these conjectures put probability one on the profile (C, c) , as Theorem 5 predicts.

It follows from the above discussion that s is a state of model $\langle T_1^M, T_2^M \rangle$ where the Aumann and Brandenburger conditions for Nash equilibrium in actions and in beliefs hold. Yet it is also true at $s = (CDDC, g_A, \beta_A, cdcd, g_B, \beta_B, m_1)$ that beliefs over decision rules are mis-specified at s , as Theorem 1 predicts, and that players are certain, and wrong, that players do not use a rational decision rule at s .

To sum it up: Theorem 6 does not say that the Aumann and Brandenburger conditions cannot be satisfied in any interim state of model $\langle T_1^M, T_2^M \rangle$. It does reveal, however, that for a large class of games, if priors are to be well-specified, so that rationality and **K'** can lead to the Aumann and Bran-

denburger's conditions for Nash equilibrium in actions,⁷ then one would have to rule out, *ex ante*, certain types of strategic uncertainty.

5.3 Similarities and differences with Nachbar (1997, 2001, 2005)

As it is quite clear, the present paper owes much of its inspiration and methodology to the work of Nachbar (1997, 2001, 2005) on repeated games. The exercise I carry out in the present paper is not, however, a straightforward translation of his research into the language of interactive epistemology and one shot games. There is much that does not translate directly, and modelling choices had to be made throughout the entire exercise. Hence I believe it is important to keep in proper perspective what the similarities and the differences are between Nachbar's research and the present paper. I explain those similarities and differences below.

For the purpose of this discussion I mostly focus the comparison to Nachbar (2005) and keep the discussion at an informal level.

(i). *The object of study.* Nachbar studies repeated games and therefore sets $\hat{\Sigma}_i$ of *behavior strategies* for player i , that is, maps from the set of finite histories to probability distributions over the action space. I study sets \hat{R}_i of *decision rules* within a context, that is, maps from an information space to the action space. The interpretation is that a decision rule specifies what a player would do given the common information that the players receive, but the rule also specifies what would have the player chosen according to that rule if the player were to receive information that is different from the one the player actually receives.

(ii) *Learnability and mutual knowledge.* Nachbar takes the belief β_i (with support given by $\hat{\Sigma}_{-i}$) of the player i over the repeated game strategy of the opponent as given and assumes that $\hat{\Sigma}_i$ satisfies learnability if for every pair of strategies from $\hat{\Sigma} := \hat{\Sigma}_1 \times \hat{\Sigma}_2$ player i (weakly) learns to predict the path of play. I say that \hat{R}_i can lead to mutual p -belief of actions (K), for a given probability level p , if for every pair of rules r_1 and r_2 from \hat{R} there is information $m_i \in M$ that allows each player i to put probability of at least p the choices made by his opponent when receiving information m_i . Notice that Nachbar's learnability assumption requires learning to take place *necessarily*

⁷Which, a fortiori, are required to hold for many games if the conditions of Theorems A or B are to hold.

for every pair of strategies in $\widehat{\Sigma}$, whereas K requires inference merely to be a *possibility* for every pair of decision rules in \widehat{R} .

(iii) *Consistency and ε -rationality.* Nachbar assumes $\widehat{\Sigma}$ satisfies ε -consistency for every player i given beliefs β_i if each player i has a uniform ε best response in $\widehat{\Sigma}_i$. I assume \widehat{R} is consistent with ε -rationality given beliefs β_i if for each player i there is a rule r_i in \widehat{R}_i such that, for every information $m \in M$, choice $r_i(m)$ maximizes expected utility for player i , where the expectation is taken with respect to the beliefs that follow from conditioning β_i on m . Theorem 1 in Nachbar (2005) finds conditions that lead $\widehat{\Sigma}$ not to be ε -consistent. Theorem 1 in the present paper finds conditions that lead \widehat{R}_i not to be consistent with ε -rationality. Hence, the structure of both theorems regarding rationality of strategies and rules is similar. What varies is the meaning of rationality in either case.

(iv) *Strategic Uncertainty.* Nachbar (2005)'s belief diversity condition (CSP) is composed of two conditions, *CS* (caution and symmetry) and *P* (pure strategies), both related to the richness of $\widehat{\Sigma}$. *CS* says that \widehat{S}_i (the set of “pure” behavior strategies in $\widehat{\Sigma}_i$) contains all pure strategies that are simple ‘translations’ of the pure strategies in \widehat{S}_{-i} and viceversa. Condition *P* says that if a behavior strategy σ_i is in $\widehat{\Sigma}_i$ then at least one pure strategy that coarsely approximates σ_i is contained in $\widehat{\Sigma}_i$ as well.

Condition **SU** in the present paper is also related to the richness of \widehat{R} . Condition **SU** is essentially like *CS* in Nachbar's (2005) paper in that it says that, given some ‘translation’ map γ between A_i and A_{-i} , if \widehat{R}_i contains r_i , then \widehat{R}_{-i} contains r_{-i} such that r_{-i} chooses $\gamma(a_i)$ at m when r_i chooses a_i at m .

A condition like Condition *P* in Nachbar (2005) is not necessary in the present paper as I only deal with “pure” decision rules here.

Because of the nature of **SU**, it is not surprising that \widehat{R}_i contains irrational decision rules. *What is surprising is that **SU**, when combined with K , makes \widehat{R}_i contain only irrational decision rules.*

(v). *Results.* Nachbar's main theorem is Theorem 1, a theorem that applies to impatient players in games that satisfy a No Weak Dominance condition and to players of any patience in games such that the pure action maxmin payoff is strictly less than the minmax payoff. Because I only deal with one shot games, in this paper discounting is not an issue and therefore I only need to assume that games satisfy the No Weak Dominance condition. It is a very weak condition.

(vi). *Interpretation.* Nachbar (2005) is quite careful in pointing out that his research does not imply that learning cannot lead to Nash equilibrium. What his research reveals is that learning will come at the cost of beliefs either being inconsistent with rationality or else certain kinds of uncertainty being ruled out at the outset. I interpret the research in the present paper in exactly the same way: one can find epistemic models with conditions that lead to Nash equilibrium, but if the beliefs in those models are to be consistent with rationality of decision rules, as defined in this paper, then they have to rule out, *ex ante*, certain types of strategic uncertainty.

5.4 Other related literature

Binmore (1987) and Canning (1992) explored the consequences of modelling rational agents by means of assuming that players use computable decision procedures in making their choices (procedures that take the decision procedure of the opponent as given) and use *the unsolvability of the halting problem* to show that there is no decision procedure that leads to rational choices against all possible decision procedures of the opponent for all games. The results presented in this paper have a similar flavor, except that they do not hinge on computability assumptions. Another difference is that, contrary to Binmore (1987) and Canning (1992), I also present the results in the language of interactive epistemology, that is, in the context of state space models where each state contains information about payoff relevant variables, as well as about the knowledge and belief by each player about each other's knowledge and beliefs.

A typical reaction to results such as those by Binmore (1987) and Canning (1992) is that they are of limited interest due to the computability assumptions made in those papers. It is argued that real decision makers would be more sophisticated than the algorithms considered by Binmore and Canning allow and that their negative results on the impossibility of players rationally understanding each other in equilibrium were not very relevant. The research presented in the present paper shows that the problems they identified are intrinsically game theoretic,⁸ can be discussed even if one abstracts from computability issues and can be readily compared with what is known about the epistemic conditions that lead to equilibrium in state space

⁸As the impossibility result in Brandenburger and Keisker (2006) is also eminently game theoretic.

models.

6 Appendix: The proof of Lemma 2

Define $w_i(a_i) := \max_{a'_i \in A_i} g_i(a'_i, \tilde{a}_{-i}(a_i)) - g_i(a_i, \tilde{a}_{-i}(a_i))$, $\underline{w}_i := \min_{a_i \in A_i} w_i(a_i)$, and $\bar{\varepsilon} := \min\{\underline{w}_1, \underline{w}_2\}$. Note that $w_i(a_i) \geq 0$ with strict inequality when a_i is not weakly dominant. Let $\underline{w}_i := \min_{a_i \in A_i} w_i(a_i) > 0$, $\bar{g}_i := \max_{a \in A} g_i(a)$ and $\underline{g}_i := \min_{a \in A} g_i(a)$. No strategy is weakly dominant, so $[\bar{g}_i - \underline{g}_i] > 0$.

For the moment fix $p \in (0, 1)$ and $\varepsilon > 0$. Pick an arbitrary $r_1 \in R_1$ and let \tilde{r}_2 be an evil twin of r_1 , namely, $\tilde{r}_2 = \tilde{a}_2(r_1(m))$ for all $m \in M$. As in the hypothesis of the Lemma, assume that $\beta_1^a(\tilde{r}_2(m); m) \geq p$ for some $m \in M$. Let $a_1^* = r_1(m)$ and $a_2^* = \tilde{r}_2(m)$ for such m .

Given such beliefs, player 1's expected payoff from choosing a_1^* is at most

$$p \cdot g_1(a_1^*, a_2^*) + (1 - p) \bar{g}_1.$$

If player 1's choice a_1^* is a best response to the player's beliefs her expected payoff from choosing a_1^* would be at least

$$p \cdot \max_{a_1 \in A_1} g_1(a_1, a_2^*) + (1 - p) \underline{g}_1.$$

Therefore, ε -rationality of r_1 requires that

$$\varepsilon + p \cdot g_1(a_1^*, a_2^*) + (1 - p) \bar{g}_1 \geq p \cdot \max_{a_1 \in A_1} g_1(a_1, a_2^*) + (1 - p) \underline{g}_1,$$

or

$$\varepsilon + (1 - p) [\bar{g}_1 - \underline{g}_1] \geq p \left[\max_{a_1 \in A_1} g_1(a_1, a_2^*) - g_1(a_1^*, a_2^*) \right]$$

but $a_2^* = \tilde{a}_2(a_1^*)$, so $\max_{a_1 \in A_1} g_1(a_1, a_2^*) - g_1(a_1^*, a_2^*) = w_1(a_1^*) \geq \underline{w}_1$. Therefore,

$$\varepsilon + (1 - p) [\bar{g}_1 - \underline{g}_1] \geq p \underline{w}_1,$$

which cannot hold, for example, for $\varepsilon < \underline{w}_1$ and $p > \frac{\varepsilon + [\bar{g}_1 - \underline{g}_1]}{\underline{w}_1 + [\bar{g}_1 - \underline{g}_1]}$. This guarantees the existence of the desired $\bar{\varepsilon}$ and \bar{p} which means, in the end, that for sufficiently low ε and sufficiently high p the rule r_1 is not an ε -rational decision rule for player 1. ■

References

- [1] Aumann, R. and A. Brandenburger, Epistemic Conditions for Nash Equilibrium, *Econometrica* 63, 1161-1180, 1995.
- [2] Binmore, K., "Modelling Rational Players, Part I," *Economics and Philosophy* 3, 179-214, 1987.
- [3] Bonanno, G., and P. Battigalli, "Recent results on belief, knowledge and the epistemic foundations of game theory", *Research in Economics* 53, 149-225, 1999.
- [4] Brandenburger, A., "Knowledge and Equilibrium in Games," *Journal of Economic Perspectives* 6, 83-101, 1992.
- [5] Brandenburger, A., "The Power of Paradox: Some Recent Developments in Interactive Epistemology," *International Journal of Game Theory*, forthcoming, 2006.
- [6] Brandenburger, A. and J. Keisler, "An Impossibility Theorem on Beliefs in Games, *Studia Logica*, forthcoming, 2006.
- [7] Camerer, C., T. Ho, and J. Chong, "A Cognitive Hierarchy Model of Games," *Quarterly Journal of Economics* 119, 861-898, 2004.
- [8] Canning, D., "Rationality, Computability, and Nash Equilibrium," *Econometrica* 60, 877-888, 1992.
- [9] Costa-Gomes, M. and V. Crawford, "Cognition and Behavior in Two-Person Guessing Games: An Experimental Study," *American Economic Review*, forthcoming, 2006.
- [10] Costa-Gomes, M., V. Crawford and B. Broseta, "Cognition and Behavior in Normal-Form Games: An Experimental Study," University of California, San Diego, Discussion Paper Series: DP 98-22, 1998.
- [11] Costa-Gomes, M., V. Crawford and B. Broseta, "Cognition and Behavior in Normal-Form Games: An Experimental Study," *Econometrica* 69, 1193-1235, 2001.
- [12] Costa-Gomes, M. and G. Weizsäcker, "Stated Beliefs and Play in Normal-Form Games," Unpublished Paper, 2005.

- [13] Dekel, E. and F. Gul, "Rationality and Knowledge in Game Theory," in Kreps, D. and K. Wallis (eds.), *Advances in economics and econometrics: Theory and Applications, vol. I*, Cambridge University Press, 1997.
- [14] Ho, T., C. Camerer and K. Weigelt, "Iterated Dominance and Iterated Best Response in Experimental 'p-Beauty Contests'," *American Economic Review* 88, 947-969, 1998.
- [15] Myerson, R., "Nash Equilibrium and the History of Economic Theory," *Journal of Economic Literature* XXXVII, 1067-1082, 1999.
- [16] Nachbar, J., "Prediction, Optimization, and Learning in Repeated Games," *Econometrica* 65, 275-309, 1997.
- [17] Nachbar, J., "Bayesian Learning in Repeated Games of Incomplete Information," *Social Choice and Welfare* 18, 303-326, 2001.
- [18] Nachbar, J., "Beliefs in Repeated Games," *Econometrica* 73, 459-480, 2005.
- [19] Nagel, R., "Unraveling in Guessing Games: An Experimental Study," *American Economic Review* 85, 1313-1326, 1995.
- [20] Samuelson, L., "Modeling Knowledge in Economic Analysis," *Journal of Economic Literature* XLII, 367-403, 2004.
- [21] Stahl, D. and P. Wilson, "Experimental Evidence on Players' Models of Other Players." *Journal of Economic Behavior and Organization* 25, 309-327, 1994.
- [22] Stahl, D. and P. Wilson, "On Players' Models of Other Players: Theory and Experimental Evidence," *Games and Economic Behavior* 10, 218-254, 1995.