Towering over Babel: Worlds Apart but Acting Together

by

Joseph Greenberg, Sudheer Gupta, and Xiao Luo *

December, 2003

Abstract

We offer a formal framework that enables the analysis of the fascinating phenomenon whereby individuals who "live in different worlds" agree to a shared course of action. We define the notion of a course of action which, unlike a (behavioral) strategy profile, does not require a complete specification of actions in every contingency. We introduce a new solution concept: a mutually acceptable course of action (MACA). Loosely speaking, an MACA consists of actions that "make sense" in every player's (perception of the) game, and every player takes into account that all players are rational. In particular, an MACA can be viewed as an (incomplete) contract or a social norm that free rational individuals would be willing to follow for their own diverse reasons. We show that by varying the degree of completeness of the underlying course of action, the concept of an MACA can be related to many of the commonly used solution concepts, such as perfect equilibrium, perfect Bayesian equilibrium, (rationalizable) self-confirming equilibrium, and rationalizable outcomes. Thus, our framework also serves to unify "game-theoretic" and "incomplete contracts" viewpoints. (JEL: C72)

*Greenberg: Department of Economics, McGill University, Montreal, Canada, H3A 2T7; Gupta: University of Michigan Business School, Ann Arbor, MI, USA, 48109; Luo: Institute of Economics, Academia Sinica, Taipei 115, Taiwan, ROC.

1 Introduction

We offer a way to formally analyze the fascinating phenomenon whereby individuals who "live in different worlds" agree to follow a common course of action. The problem of analyzing shared action among individuals has attracted the attention of sociologists, psychologists, political scientists, and organizational theorists for a long time. While there are differences in focus between different research streams, the underlying problem is the same: *why, and when, do rational actors choose to follow a common course of action*? Central to this question is the notion of differences in actors' perceptions of the situation; their values, beliefs, and views of the world. There is ample evidence that several political and military situations are best analyzed as different games viewed by different players. For example, Snyder (1979) and Boudon (1981) argue that during the first world war, Germany believed that the game being played was a prisoner's dilemma type game while Britain believed that the game being played was a coordination type game. Other examples include the German occupation of France in 1940, the Allied invasion of Normandy in 1944, and the Middle East war of 1973 (see Section 5.1.2 for a discussion). Finally, the fact that resources are devoted to diversion and deception[1] demonstrates that countries, politicians, military generals, and other economic agents realize that different players may (be induced to) analyze different games.

And, there is a rich tradition of the idea of "different worlds" in literature, a few examples here being: Moliere's and Shakespeare's "comedies of errors" where the identities of "players" are misperceived; Akatagawa's (1952) "In a wood" (or, its well-known movie version, "Rashomon", by Kurosawa) where players' perceptions of an event they witness differ drastically; and Canetti's (1946) "Auto-Da-Fe",[2] where there is no ambiguity concerning the perceptions of the ("physical") identities or actions that are actually being taken, but individuals attribute very different

[1]See Crawford (2003) and Hendricks and McAfee (2003) for several such examples of feints, or 'strategic misrepresentation', from oil drilling to ex-President George Bush's famous 1988 campaign promise: "Read my lips: no new taxes".

[2]The novel was published in German as "Die Blendung" (The Deception) in 1935; the first US version appeared under the title "The Tower of Babel" in 1947.

motivation and interpretations to the actions the other players take.

At first glance, the notion that different individuals may perceive a given situation as different games may seem like a radical change from "classical" game theory. However, it turns out that our framework facilitates formal representation of "mainstream" arguments as well, including those concerning the most fundamental notions in game theory: the Minmax theorem and the Nash equilibrium. Indeed, von Neumann and Morgenstern themselves used the idea of worlds apart to justify the minmax value in a two-person zero-sum game. They associated with such a game two auxiliary games: the majorant game and the minorant game, which correspond to each player assuming that he moves first, and his choice is observed by the other player. Thus, each player is analyzing a different (perfect information extensive form) game, where each player sees himself as the "leader" and the other player as the "follower" (see Section 5 for a formal discussion).

> *Indeed, the rules of the game prescribe that each player must make his choice (his personal move) in ignorance of the outcome of the choice of his adversary. It is nevertheless conceivable that one of the players, say 2, "finds out" his adversary; i.e., that he has somehow acquired the knowledge as to what his adversary's strategy is. The basis for this knowledge does not concern us; it may (but need not) be experience from previous plays. At any rate we assume that the player 2 possesses this knowledge. ... Under these assumptions we may then say that player 2 has "found out" his adversary. [von Neumann and Morgenstern, 1947, p. 105.]*

More recently, Aumann and Brandenburger (1995) interpreted a Nash equilibrium as players' choices when each player knows his own payoff functions and strategy choices of the other players. As they write (*ibid*, p. 1161): "Suppose that each player is rational, knows his own payoff function, and knows the strategy choices of the others. Then the players' choices constitute a Nash equilibrium in the game being played." As in von Neumann and Morgenstern's interpretation of the minmax value, this argument also implicitly suggests that different players perceive and analyze different games.[1]

There are, however, several fundamental difficulties in formally applying traditional game-theoretic analysis to situations with divergent perceptions, as such analysis almost always assumes that the players play the same game. That is, following Harsanyi (1967-68), each player is assumed to be (at least probabilistically) aware of all the available actions and payoffs of all the other players, and this information as well as the structure of the game is common knowledge. In particular, all players consider the same "states of Nature". This assumption prevents the analysis of situations where different players may perceive different games.

Even when the game is common knowledge, there are two additional limitations of most of the existing models in game theory that we seek to relax. First, most equilibrium solution concepts usually employed in applications impose (different degrees of) a stringent requirement: commonality of beliefs among players regarding other players' actions, including actions in the "off-equilibrium" contingencies (that are not supposed to arise if the equilibrium contract is followed). Second, in the most commonly used equilibrium notions (such as perfect or sequential equilibrium) players are required, in addition to agreeing on the precise actions to be taken in every contingency (whether or not it arises), to also agree on the precise way in which players might deviate (or commit mistakes - known as "trembles") from these specified actions. Many game theorists have pointed out the problems with these assumptions (e.g., see Kreps, 1990; Rubinstein, 1991).

In sum, our analysis of shared action with divergent perceptions enables us to extend the literature by offering a framework for studying the following situations: a) when players perceive different games; b) when the game is common knowledge but players' beliefs may differ off the equilibrium path; and, c) when the game is common knowledge and the actions to be taken in every contingency are agreed upon, but players may not necessarily have the same beliefs on how other players might "tremble" (or make mistakes).

[1] To be sure, no formal framework is presented by these authors that could accommodate players playing different games.

We carry out our analysis using the notion of a "*course of action (CA)*" which, unlike a strategy profile, does not require a complete contingent specification of actions in all information sets. Informally, a course of action is a specification of (possibly probabilistic, or mixed) actions to be taken in some, but not necessarily all, contingencies. A course of action can be interpreted as a social norm or an incomplete contract. A contract (or agreement) may be incomplete in the sense that it does not specify actions in all possible contingencies, and it may be partial in the sense that players may agree to take certain actions towards resolution of a situation, without actually fully resolving it. Such partial and incomplete contracts form a large part of everyday human interactions that form the basis of many social and economic phenomena. However, the solution concept ("equilibrium") employed in most game-theoretic analyses is a strategy profile, which, by definition, provides a complete specification of actions in all possible contingencies. Thus, the existing models in game theory are not easily amenable to the analysis of incomplete contracts.

Our solution concept, a *mutually acceptable course of action (MACA)* is, loosely speaking, a course of action that "makes sense" in every player's (perception of the) game. That is, a course of action is *mutually acceptable* if no player would wish, in his own world, to deviate from it. When deciding on whether or not to deviate from a course of action, every player takes into account that all players are "rational". In making their decisions, each player analyzes possible consequences of deviations from the proposed course of action. Players would be willing to conform to a proposed course of action as long as their conformity does not conflict with rational behavior. Observe that each player may rationalize his expectations in a different way, as long as this does not violate the common knowledge of rationality as perceived by each player.

In the "classical" case where all players play the same game, and this is common knowledge, we relate our solution concept to several equilibrium notions frequently used in game-theoretic analysis and applications, by exploring the implications of different degrees of completeness of the mutually acceptable course of action. Specifically, we show that when the underlying course of action is "complete" (i.e., specifies an action in all possible contingencies), the MACA lies "in-between" perfect equilibrium (Selten, 1971) and perfect Bayesian equilibrium (Fudenberg and Tirole, 1991). This is true even in the special cases of normal form[3] games, and in games with perfect information.

If the course of action specifies a "path", that is, leaves out contingencies if someone deviates from the specified actions, the MACA refines self-confirming (Fudenberg and Levine, 1993) and rationalizable self-confirming equilibria (Dekel, Fudenberg, and Levine, 1998, 2002).

If the course of action is completely silent, that is, does not specify any actions at any nodes, the MACA refines rationalizable outcomes in normal form games (Bernheim, 1984; Pearce, 1984). For extensive form games, the notion of a null MACA provides a new and attractive (see Example 4.3.2) definition of rationalizability.

As one application of our general framework, we can interpret our results in the language of contract theory. As noted earlier, a course of action can be seen as a proposed incomplete contract that does not specify parties' obligations in all contingencies. An MACA can be viewed as an incomplete contract that rational individuals would be willing to follow for their own (possibly diverse) reasons. As a large body of literature in contract theory has argued, rarely, if ever, does a contract specify actions in all possible contingencies (e.g., see Williamson, 1985; Grossman and Hart, 1986; Hart and Moore, 1988). Most of this literature sees contractual incompleteness as resulting from "unforeseen contingencies", bounded rationality, or certain "transaction costs" that make it impossible to have a complete contingent contract describing the terms in each possible state of nature. Thus, there is an implication that loss of contractual completeness causes some loss of efficiency, that is, all parties could be better off if a complete contract were possible. One exception to this literature is a recent paper by Bernheim and Whinston (1998) who show that, in presence of non-verifiable parameters, making a contract more incomplete can lead to better outcomes even in the absence of any costs of writing (and enforcing) contracts. We obtain a similar (and stronger) result, that an incomplete MACA may Pareto-dominate any *complete* MACA. However, our motivation and framework are very different (see Example 2.1 and the discussion there).

[3]Some readers may prefer the term 'simultaneous-move' games.

We will formalize these ideas and demonstrate their usefulness in the rest of the paper. The paper is organized as follows. For ease of exposition, and to facilitate comparison with existing game-theoretic notions, we present most of our analysis in the classical case where all players face the same game which is assumed to be common knowledge. We start in Section 2 with two examples to illustrate the main ideas. Section 3 lays out the basic notation and the key definitions of a CA and an MACA, and examines some basic properties of an MACA. Section 4 shows how an MACA is related to some of the commonly used equilibrium concepts (definitions of these concepts are summarized in Appendix 1). Section 5 extends this analysis to the case where players live in different worlds. Appendix 2 contains all the proofs.

2 Examples

To demonstrate the main ideas, we present two simple examples before a formal analysis. Here, we restrict ourselves to the classical case where the game is common knowledge. The first example illustrates that an incomplete contract can Pareto dominate every complete contract even in a game of perfect information.

Example 2.1: Consider the perfect information game depicted in Figure 2.1. Player 1 (an entrepreneur) has to decide whether to sell the firm (action L_1) or to delegate control to his son (player 2). Player 2 can then decide to manage the firm himself (action L_2) or hire player 3 (CEO) to run the business. The CEO, in turn, may or may not delegate control to player 4 (a subordinate, a manager, or a team of managers represented as a single player). The manager can, then, either exert effort to manage the business well (action L_4), or shirk (action R_4). Assume that the game, actions, and resulting payoffs as depicted in Figure 2.1 are all common knowledge.[4]

[4]Please note that payoffs here are given in utils, so any penalties for the manager for shirking are accounted for. Our interest here is in partial contracts rather than incentives.

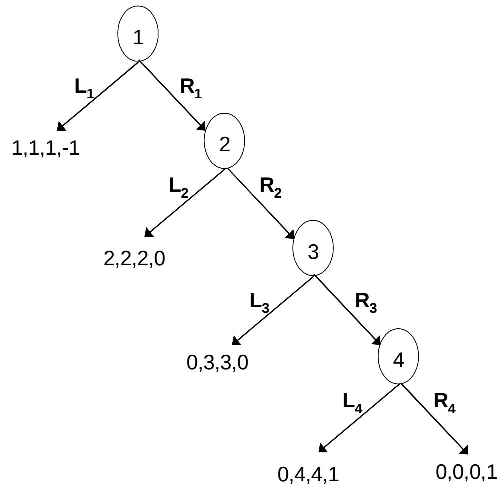

Figure 2.1

There are two subgame perfect equilibria of this game: (L_1, R_2, R_3, L_4) and (L_1, R_2, L_3, R_4). Both of them support the action L_1whereby the entrepreneur sells the firm. The unique subgame perfect equilibrium payoff is (1,1,1,-1).

However, we shall now argue that players can do better through an incomplete contract that does not specify the actions players 3 and 4 are to take at their respective information sets. Consider the (incomplete) course of action $x = (R_1, L_2)$ (i.e., the CA specifies that the entrepreneur would delegate control to his son, who chooses not to delegate further). Note that x is not a strategy profile as it does not specify an action for each player in all possible contingencies. We show (somewhat informally, for now) that x constitutes a mutually acceptable course of action (MACA); that is, all players will rationally choose to abide by this course of action. Observe that this CA generates the payoff (2,2,2,0), which Pareto dominates the unique subgame perfect equilibrium payoff (1,1,1,-1).

The reasoning is as follows. If player 4's information set is reached, it is rational for player 4 to choose either L_4 or R_4 as he is indifferent between the two. Since the CA (R_1, L_2) does not specify an action for player 4, player 3 can rationally choose L_3 (if he believes that 4 will shirk and choose R_4) or R_3 (if he believes 4 will work hard). Thus, player 2, contemplating a deviation from the CA (R_1, L_2), can reasonably expect that 3 might choose R_3 and that 4 might choose R_4. Therefore, player 2 will be willing to follow x and hence player 1 will also be willing to follow x. Note that it is only because the CA x is incomplete that it is mutually acceptable; no course of action that specifies "rational actions" for the players in all information sets can support the CA x. In particular, no subgame perfect equilibrium supports the path (R_1, L_2) because any such equilibrium requires that all players have the same beliefs regarding other players' actions (including those at nodes that are precluded from being reached if the equilibrium is followed). In contrast, (R_1, L_2) is an MACA because players 2 and 3 need not have (or, may not be aware that both, do in fact have) the same beliefs regarding player 4's behavior if he is delegated control of the firm.

This example highlights the important observation that the requirement that all players have the same beliefs regarding other players' behavior is constraining and may preclude reasonable, in fact, more plausible, predictions in many applications. Observe that incomplete contracts arise here even in absence of indescribable contingencies and transaction costs.

This example also illustrates that incomplete contracts may lead to better outcomes than those achievable through complete contracts. That is, ambiguity about actions at some information sets may yield Pareto-dominating payoffs. This is in sharp contrast to most of the literature on contract theory, a notable exception being the recent paper by Bernheim and Whinston (1998). They show that a less complete contract may Pareto dominate a more complete contract; thus, more contractual ambiguity may be beneficial, in agreement with our argument above. However, in their model (which differs from ours) a "complete" contract cannot be Pareto improved. The optimality of incomplete contracts in their setting is conditional on the presence of non-verifiable parameters that cannot be specified in any contract.

The next example illustrates how incomplete contracts can arise in a game of imperfect information. The basic insight from Example 2.1 carries over: some (or all) of the players may benefit by leaving their choice of actions ambiguous in some information sets.

Example 2.2: [5]Consider the imperfect information game depicted in Figure 2.2, which may represent the following diplomatic "peace-negotiation" scenario. Each of the two warring countries, 1 and 2, has to decide whether or not to reach a peace agreement, represented by the path (R_1, R_2). If the countries fail to reach an agreement, country 3 would "re-evaluate" its policy, a decision that would affect both countries 1 and 2. Assume that country 3 has no way to know which of the two countries caused the breakup of the negotiations (otherwise, it could threaten to retaliate against that country). All it observes is whether or not the negotiations were successful. As the payoffs in Figure 2.2 indicate, it is in the best interest of country 3 that the two warring countries sign the peace agreement. This game has a unique (mixed strategy) Nash equilibrium given by: Player 1 uses the mixed strategy $(\frac{1}{2}\ L_1,\ \frac{1}{2}\ R_1)$, player 2 uses the pure strategy L_2, and player 3 uses the mixed strategy $(\frac{1}{2}\ L_3,\ \frac{1}{2}\ R_3)$. Player 3's Nash equilibrium payoff is, therefore, 1/2.

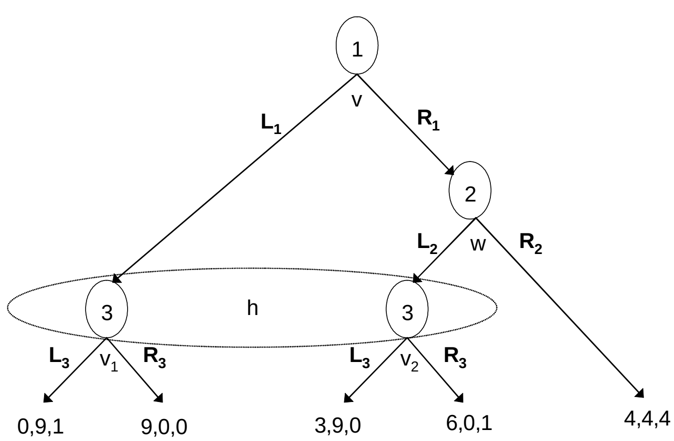

Figure 2.2

We shall now argue that player 3 may expect a payoff of 4 by not specifying the action he might take (if information set h were to be reached). That is, the (incomplete) agreement

[5]This example also appears in Greenberg (2000).

$x = (R_1, R_2)$ is an MACA: both players 1 and 2 will choose to follow this course of action. Indeed, player 1 as well as player 2 may fear that in the event of a deviation from x (i.e., breakdown of negotiations), player 3 will retaliate against him.[6] Thus, the path (R_1, R_2) will be followed because x does not specify player 3's action. This outcome could not be supported in any contract that specified player 3's action, or alternatively, if players 1 and 2 shared the same beliefs about 3's action and this was common knowledge.

We now formalize these ideas in the next sections.

3 MACA in a Common World

We start with the classical case where all players face the same game (possibly with imperfect information), which is assumed to be with perfect recall and common knowledge. We denote the game by $T = (N, V, H, \{A(h)\}_{h \in H}, \{u^i\}_{i \in N})$, where $N = \{1, 2, \ldots, n\}$ is the set of players, V is the set of nodes (or vertices), H is the set of information sets, $A(h)$ is the set of pure actions available at information set h, and u^i is player i's payoff function. For simplicity, we shall assume that the set of vertices is finite.

A mixed action at information set h is a probability distribution over the pure actions in $A(h)$. We denote the set of mixed actions at h by $\Delta(h)$, and the set of information sets belonging to player i by H^i. A behavioral strategy for player i is a function that assigns to every $h \in H^i$ a mixed action from $\Delta(h)$.[7] The set of strategies of player i is denoted by $\mathbb{Y}^i$, and the set of strategy profiles is denoted by $\mathbb{Y}$ (i.e., $\mathbb{Y} = \times_{i \in N} \mathbb{Y}^i$). For $y \in \mathbb{Y}$, we denote by $y(h)$ the mixed action of y at h and by $y(-h)$ the profile of mixed actions of y at all information sets other than h. For $y \in \mathbb{Y}$, $u^i(y)$ denotes i's (expected) payoff if y is followed from the root of the game.

[6]Player 1 will not deviate from x and choose action L_1 since he might fear that player 3 might believe that he is at vertex v_1 and therefore may choose action L_3, yielding player 1 the payoff of 0. Similarly, player 2 will not deviate from x and choose action L_2 fearing that player 3 might believe that he is at vertex v_2 and therefore may choose action R_3, yielding player 2 the payoff of 0.

[7]Throughout this paper, we will use the term strategy to mean behavioral strategy.

We can now formalize the basic building blocks in our analysis: the notion of a course of action and the conditions under which it will be mutually acceptable to all players.

Definition 3.1 *A course of action (CA) is a mapping* $x : H \to \cup_{h \in H} \Delta(h) \cup \{\emptyset\}$*, with* $x(h) \in \Delta(h) \cup \{\emptyset\}$ *for all* $h \in H$*.*

Thus, a course of action may or may not specify a mixed action at an information set. The interpretation of $x(h) = \emptyset$ is that the CA x does not specify which action from $\Delta(h)$ player i would take at h, where $h \in H^i$; otherwise, $x(h)$ specifies player i's (mixed) action at h. In particular, a CA x is said to be *complete* if $x(h) \neq \emptyset$ for all $h \in H$. A complete CA is therefore a strategy profile.

The central question with which we are concerned here is: *which course of action would be followed by free rational individuals?* Keeping with the noncooperative approach, a course of action will be followed if no single player has an incentive to deviate from it. Such a course of action will be called "mutually acceptable" (we abbreviate it as MACA). Informally, an MACA is a CA x such that following x is always optimal for every player, who is aware that, like himself, all other players are "rational". We first elaborate on the criterion of rationality.

As is customary, we impose the rationality criterion through the requirement that actions chosen by a player be best responses to his beliefs about opponents' actions. In order to ensure that players' actions are optimal in *every* contingency, including contingencies that might not actually arise, we follow Selten's (1975) idea of "trembles". Specifically, each player i associates with strategy $y^j \in \mathbb{Y}^j$ of every player $j \neq i$ a "trembling sequence" $\left\{y_k^j\right\}_{k=1}^{\infty}$ of totally mixed strategies[8] that converges to y^j. Let $y_k^j \rightsquigarrow y^j$ denote such a sequence. A strategy $y^i \in \mathbb{Y}^i$ is player i's perfect best response to $y^{-i} \in \mathbb{Y}^{-i}$ if actions specified by y^i remain optimal for i along the trembling sequences.

In most cases, however, player i does not know (or does not have a degenerate point-expectation about) the precise strategy y^j that player j might adopt. Rather, i is likely to

[8] A *totally mixed* strategy for a player assigns strictly positive probability to every action at every information set of that player.

have a set of strategies, $Y^j \subseteq \mathbb{Y}^j$ that i believes j might adopt. In that case, i assigns some probability distribution over Y^j. By Claim 0 in Appendix 2, for any distribution over Y^j, there exists an outcome equivalent distribution over Y^j that has a finite support (i.e., these two distributions yield, for any $y^{-j} \in \mathbb{Y}^{-j}$, the same distribution over the terminal nodes). That is, player i believes that player j will choose $y^j_t \in Y^j$ with probability λ_t, $t = 1, 2, ..., m$. Recall that perfection requires that player i associates y^j_t with some sequence $y^j_{t,k} \rightsquigarrow y^j_t$. Thus, λ is uniquely associated with a sequence $\left\{y^j_k\right\}_{k=1}^{\infty}$ of totally mixed strategies of player j such that the k^{th} element, y^j_k, along this sequence is outcome-equivalent to choosing $y^j_{t,k}$ with probability λ_t, $t = 1, 2, ..., m$. We denote by $y^j_k \overset{Y^j}{\rightsquigarrow} y^j$ such a sequence with a limit of y^j. (Observe that when $Y^j = \{y^j\}$ is a singleton, we have that $y^j_k \overset{Y^j}{\rightsquigarrow} y^j$ if and only if $y^j_k \rightsquigarrow y^j$.) Our rationality criterion requires that every action chosen by a player be optimal along these "trembling" sequences. More specifically, we require that for a strategy $y^i \in Y^i$ of player i, there exist $y^i_k \rightsquigarrow y^i$ and $y^j_k \overset{Y^j}{\rightsquigarrow} y^j$ for all $j \neq i$ such that, for all $h \in H^i$, and for all $k = 1, 2, \ldots,$ $u^i(y(h), y_k(-h)) \geq u^i(a, y_k(-h))$ for all $a \in \Delta(h)$.

We can now define our main solution concept: a "mutually acceptable course of action". As noted earlier, a mutually acceptable course of action is a course of action that rational individuals agree to follow for their own, possibly different reasons. Players would be willing to conform to a proposed course of action as long as their conformity does not conflict with rational behavior.

Definition 3.2 *A CA x is a mutually acceptable course of action (MACA) if there exists a set of strategy profiles $Y \equiv Y^1 \times Y^2 \cdots \times Y^n$ that supports x. That is, for every player i and every $y^i \in Y^i$ there exist $y^i_k \rightsquigarrow y^i$ and $y^j_k \overset{Y^j}{\rightsquigarrow} y^j$ for all $j \neq i$ such that*

(i) for all $h \in H$, $y(h) = x(h)$ whenever $x(h) \neq \emptyset$, and

(ii) for all $h \in H^i$ and for all $k = 1, 2, \ldots,$ $u^i(y(h), y_k(-h)) \geq u^i(a, y_k(-h))$ for all $a \in \Delta(h)$.

Remarks: (1) Observe that the chosen sequence $\{y_k\}_{k=1}^{\infty}$ in Definition 3.2 depends on player

i, because each player can rationalize his choice of $y^i \in Y^i$ using his beliefs, y^{-i}, about the strategies that the other players might employ. When needed, we shall emphasize this fact by denoting this sequence as $\{y_k[i]\}_{k=1}^{\infty}$.

(2) Definition 3.2 requires that all players "support" an MACA x by the same set of strategy profiles, $Y \subseteq \mathbb{Y}$. One could justify this assumption by the fact that as the game T and the rationality of the players are common knowledge, every inference player i can make about player j's plausible choices, can also be made by any other player. However, given our focus on divergent beliefs, we stress here that the above definition would not change even if we allow different players to support x with different sets of strategy profiles.[9]

Before we explore some special cases of interest for our solution concept and its relation to existing game-theoretic notions, we present two important properties that will be useful in understanding and establishing our results. First, an intuitively appealing property of an MACA is that it is not supported by strictly dominated strategies. Moreover, in the case of normal form games, an MACA is not supported by weakly dominated strategies.

Claim 3.3 *Suppose Y supports an MACA in a game T. Then, for each player i, $y^i \in Y^i$ is not a strictly dominated strategy. Moreover, if T is a game where each player has only one information set (in particular, if T represents a normal form game), then $y^i \in Y^i$ is not a weakly dominated strategy.*

The second property below states that acceptability cannot be lost by making the course of action less complete. That is, the more stringent a contract is, the less likely it is to be accepted. It is *ambiguity* that gives people hope; shared action is made possible by lack of specification of actions that one or more players are likely to take at some information sets in the game, as shown in examples 2.1-2.2. Formally,

[9] To see this, observe that if for every $j \in N$ there exits a set $Y[j]$ that supports x, then the (single) set $Y \equiv \times_{i\in N} \cup_{j\in N} Y^i[j]$ also supports x. This is a special case of a more general result (Theorem 3.4) in Greenberg, Monderer, and Shitovitz (1996).

Claim 3.4 *Let x be an MACA, and let y be a CA such that $y(h) = x(h)$ whenever $y(h) \neq \emptyset$. Then, y is an MACA.*

4 Applications in a Common World

We now explore the concept of a course of action in greater detail by looking at a few special cases. Our objective here is to demonstrate the generality of the notion of an MACA, its relationship to existing equilibrium notions, and its potential usefulness in a wide range of applications. We show in this section that a number of different equilibria often employed in applications can be better understood as variants of our equilibrium notion, depending on the degree of completeness of the acceptable course of action. Recall that an MACA may range from specifying an action in all contingencies, to being completely silent, representing different degrees of contractual completeness. In particular, an MACA need not even specify an entire path. Indeed, we often witness partial agreements where players agree to take some actions towards a resolution, without agreeing on the "final outcome".[10] Similarly, contracts may be incomplete to different degrees: some may specify consequences for most, but not all, deviations from the proposed equilibrium, while others may not specify any such consequences. Our framework can easily accommodate such variations. We restrict our analysis here to three special cases:

(i) a complete MACA — a course of action that specifies actions in all possible contingencies; this represents a complete contingent contract or agreement. A complete CA is a strategy profile.

(ii) a path MACA — a course of action that specifies an action at the root of the game and at every information set that is reached with positive probability if the CA is followed. This is a typical representation of an incomplete contract where no contingencies are covered if

[10]This is a common feature of most political agreements, like the Dayton accord, the Israeli-Palestenian accord of 1996, arms treaties like Anti-Ballistic Missile treaty. At most, participants in such treaties hope for some actions towards an intended outcome, but the 'terminal nodes' are far from certain. The same is true in most ongoing (business or personal) relationships.

someone deviates from the specified contract.

(iii) a null MACA — which does not pin down any actions at any information set; players can only rely on deductions they can make from the common knowledge of the game and rationality of all the players to anticipate how the game will play out.

4.1 Complete MACA

A course of action x is *complete* if $x(h) \neq \emptyset$ for all $h \in H$. That is, players specify an action in every possible contingency they may face. This assumption is the building block of most game-theoretic analysis, based on the notion of a strategy profile, which, as we have argued earlier, is quite restrictive. Nevertheless, we proceed in this subsection with the assumption of a complete course of action, while retaining our focus on divergent perceptions within the same game that is common knowledge. We will show that even though players agree on a "complete contract", they may do so for their own reasons (or beliefs). Since we believe this to be a widely prevalent characteristic of human interactions, our framework and solution concept may provide a better way of modeling some situations.

Our first result asserts that a complete contract is an MACA if and only if it is "self supporting". However, while all players believe that the complete MACA will be followed, it may be impossible for the players to believe in the same sequence of trembles that converge to this MACA. (See Example 4.1.4 below.)

Claim 4.1.1 *The CA x is a complete MACA if and only if it is supported by the set $Y \equiv \{x\}$.*

The following Claim offers an alternative definition for a complete MACA, and also illustrates the fact that we allow for divergent beliefs.

Claim 4.1.2 *A strategy profile x is a complete MACA if and only if for each player i there exists a sequence $y_k[i] \rightsquigarrow x$ such that, for all $h \in H^i$ and for $k = 1, 2, \ldots,$*

$$u^i(x(h), y_k[i](-h)) \geq u^i(a, y_k[i](-h)) \text{ for all } a \in \Delta(h).$$

Most applications of dynamic games, such as mechanism design and signaling games, use some variants of perfect equilibrium (Selten, 1975). We shall show that the set of complete MACA is "in-between" the sets of perfect equilibria and of perfect Bayesian equilibria (Fudenberg and Tirole, 1991).[11] As every totally mixed Bayesian perfect equilibrium is a perfect equilibrium, it follows that the only difference between these three notions (perfect, Bayesian perfect, and complete MACA) lies in the way players may reason in information sets that are not reached. We will elaborate on these differences below.

Our first result establishes the formal relationship between a complete MACA and perfect equilibrium.

Claim 4.1.3 *Every perfect equilibrium is a complete MACA.*

Example 4.1.4 below demonstrates that, in general, a complete MACA need not be a perfect equilibrium. This is because of the more stringent requirement imposed by the definition of a perfect equilibrium that all the players must have the same beliefs at all information sets not only concerning the actions that will be taken there, but also on the way these information sets were reached. While commonality of beliefs along the equilibrium path can be justified (e.g., "past observations"), this requirement is not realistic for information sets off the equilibrium path, as noted by others as well (e.g., see Kreps 1990, p.166). That is, for shared action to take place, players have to agree on what that action is going to be, but not necessarily on the trembles that resulted in an information set being reached off the equilibrium path.

Example 4.1.4: Consider the game depicted in Figure 4.1.4.

[11] See Appendix 1 for formal definitions of these concepts.

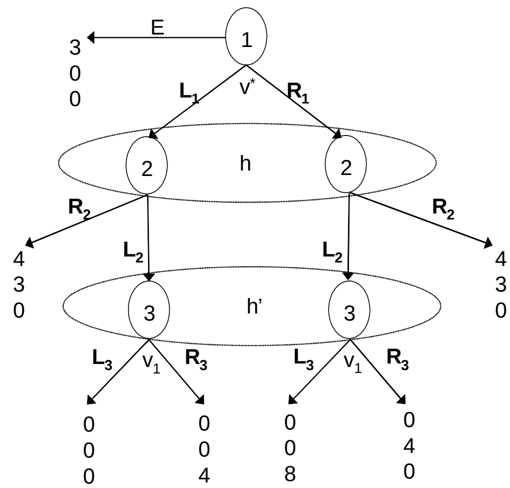

Figure 4.1.4

Define the complete CA x by: $x(v^*) = E$, $x(h) = L_2$, and $x(h') = R_3$. We shall now show that x is an MACA but it is not a perfect equilibrium. It is easy to verify that the condition in Claim 4.1.2 is satisfied for the following sequences $\{y_k[i]\}_{k=1}^{\infty}$ and, hence, x is an MACA. For player i, $i = 1, 2, 3$, define $y_k[i]$ as follows:[12]

$$
\begin{array}{lll}
y_k[1](v^*) = \left(1 - \frac{2}{k}\right) E + \frac{1}{k} L_1 + \frac{1}{k} R_1 & y_k[1](h) = \left(1 - \frac{1}{k}\right) L_2 + \frac{1}{k} R_2 & y_k[1](h') = \frac{1}{k} L_3 + \left(1 - \frac{1}{k}\right) R_3 \\
y_k[2](v^*) = \left(1 - \frac{5}{k}\right) E + \frac{1}{k} L_1 + \frac{4}{k} R_1 & y_k[2](h) = \left(1 - \frac{1}{k}\right) L_2 + \frac{1}{k} R_2 & y_k[2](h') = \frac{1}{k} L_3 + \left(1 - \frac{1}{k}\right) R_3 \\
y_k[3](v^*) = \left(1 - \frac{5}{k}\right) E + \frac{4}{k} L_1 + \frac{1}{k} R_1 & y_k[3](h) = \left(1 - \frac{1}{k}\right) L_2 + \frac{1}{k} R_2 & y_k[3](h') = \frac{1}{k} L_3 + \left(1 - \frac{1}{k}\right) R_3
\end{array}
$$

Observe that the sequence $\{y_k[2]\}_{k=1}^{\infty}$ implies that player 2 believes that if player 1 deviates from E, then it is more likely that player 1 would choose R_1. Sequence $\{y_k[3]\}_{k=1}^{\infty}$ implies that player 3 believes that if player 1 deviates from E, then it is more likely that player 1 would choose L_1. But perfect (or sequential) equilibrium does not allow for such divergent beliefs about trembles; all players must hold the same beliefs about the way an information set was reached. In our example, for player 2 to choose L_2 it must be the case that 2 believes that

[12] Here, for example, $(1 - 1/k)\, L_2 + (1/k) R_2$ means that player 2 chooses action L_2 with probability $(1 - 1/k)$ and action R_2 with probability $1/k$.

player 1 would tremble in such a way that the probability of R_1 is not lower than 3/4, but with these same beliefs player 3 would be better off choosing L_3. Thus, x is an MACA but it is not a perfect (or sequential) equilibrium.

Note that the MACA x in the above example is not totally mixed. (The probability with which player 1 will use either L_1 or R_1 is 0, and it is this fact that allows the players to hold different beliefs about the trembles, even though the MACA is complete.) For a totally mixed complete MACA, the following holds:

Claim 4.1.5 *A totally mixed complete CA is an MACA if and only if it is a perfect equilibrium.*

Next, we relate a complete MACA to perfect Bayesian equilibrium.

Claim 4.1.6 *Every complete MACA is a perfect Bayesian equilibrium.*[13]

Claims 4.1.3 and 4.1.6 together imply that a complete MACA lies "in-between" perfect and perfect Bayesian equilibrium. In general, the set of perfect equilibria may be a strict subset of the set of complete MACAs which, in turn, may be a strict subset of the set of perfect Bayesian equilibria. Example 4.1.4 demonstrates the validity of the first part of this assertion. To see the second part, consider the case of normal form games. Note that a perfect Bayesian equilibrium coincides with a Nash equilibrium; hence, it may involve weakly dominated strategies. However, by Claim 3.3, a complete MACA never employs weakly dominated strategies in normal form games. Thus, a complete MACA provides a useful refinement of perfect Bayesian equilibrium.

The same is true for games with perfect information: the set of complete MACAs may be a strict subset of the set of subgame perfect equilibria and may strictly include the set of Perfect equilibria. The weak inclusions are implied by Claims 4.1.2 and 4.1.6. To realize the possibility of the strict inclusions, consider the following two examples.

Example 4.1.7: In the game depicted in Figure 4.1.7, the strategy (r_1, r_2) is a subgame perfect equilibrium. However, the strategy r_1 is weakly dominated, for player 1, by the strategy l_1.

[13]See Definition A1.2; the result would hold even if we add the requirement in the definition of PBE that Bayes' rule is used 'wherever possible'.

Hence, by Claim 3.3, the action r_1 can never be part of an MACA. Indeed, if player 1 assigns any trembles to player 2, his unique best response is l_1. Therefore, (r_1, r_2) is not a complete MACA. The only complete MACA in this game coincides with the unique perfect equilibrium: (l_1, r_2).

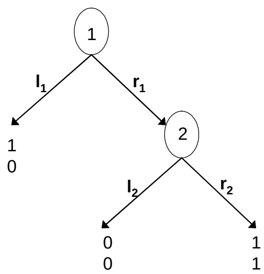

Figure 4.1.7

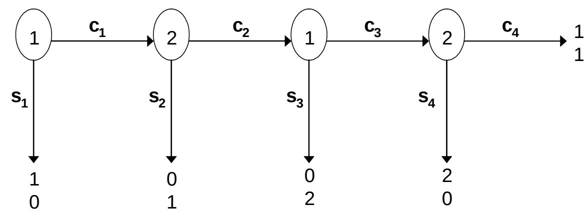

Figure 4.1.8

Example 4.1.8: The game depicted in Figure 4.1.8 shows that a complete MACA need not be a perfect equilibrium. Indeed, consider the strategy profile $y = (c_1, c_2, c_3, c_4)$ in this game. We claim that y is not a perfect equilibrium, but it is a complete MACA. To see the first point, assume in negation that y is a perfect equilibrium, supported by a sequence $\{y_k\}$ of totally mixed strategies. Let $p_k(a)$ denote the probability of playing an action a under y_k. For c_2 to be player 2's local best response to y_k it must be the case that $p_k(s_3) \geq p_k(s_4)$. But then, as $p_k(s_2) > 0$, it follows that player 1's unique local best response to y_k at the root of the game is action s_1, and not c_1. Thus, y cannot be a perfect equilibrium.

To realize that y is a complete MACA, by Claim 4.1.2 we need to show that for each player i, there exists a sequence $\{y_k[i]\}$ such that i's actions in $y = (c_1, c_2, c_3, c_4)$ are his local best response to this sequence. For player 1, let $\{y_k[1]\}$ be such that $p_k(s_2) + p_k(s_3) \leq p_k(s_4)$, and for player 2 let $\{y_k[2]\}$ be such that $p_k(s_3) \geq p_k(s_4)$. Then, $\{y_k[1]\}$ and $\{y_k[2]\}$ satisfy the inequality in Claim 4.1.2; therefore, y is a complete MACA.

4.2 Path MACA

A path MACA is a CA x that specifies a (mixed) action at the root of the game and at every information set that is reached with positive probability if x is followed. That is, unlike a strategy profile, a path MACA does not specify actions in all contingencies.

In general, a path MACA need not be derived from a Nash equilibrium, as was illustrated in Example 2.2. Nor is the converse true, not even in normal form games. To realize this assertion, note that a path MACA in a normal form game is a complete MACA. Let x be a Nash equilibrium in a normal form game that involves weakly dominated strategies. Assume, in negation, that x is a path MACA. By Claim 4.1.1, it is supported by $Y \equiv \{x\}$. But then, by Claim 3.3, x cannot involve weakly dominated strategies. Thus, x cannot be a path MACA.

However, as we will now show, every path MACA that satisfies the "unique deviator" property can be derived from a Nash equilibrium. For a CA x, denote by $H(x)$ the set of information sets that are reached (with positive probability) if x is followed.[14]

Definition 4.2.1 *(Unique Deviator Property) A path CA, x, has the unique deviator property if for each $h \notin H(x)$ there is a unique player $i(h)$ such that no deviation from x by any other player $j \neq i(h)$ can result in reaching h. That is, if $h \notin H(x)$, then $h \notin H(y)$ for any strategy y such that $y(h') = x(h')$ for all $h' \in H(x) \cap H^{i(h)}$.*

Note that every path CA x in a game with perfect information, as well as in any normal form game, has the unique deviator property. Games with this property (i.e., where every path in the game has a unique deviator) are called games with *observed deviator* property (Fudenberg and Levine, 1993). To realize the difference between these notions, consider an extensive form game where a player has two choices at the root of the game: action L ends the game, and action R leads to a proper subgame T' which does not have the observed deviator property. Clearly, T does not have the observed deviator property, but the path $\{L\}$ does have the unique deviator

[14] An information set h is said to be reachable from an information set h' if there is a vertex v in h such that the path connecting the root of the game with v passes through a vertex w in h'.

property. We then have the following result.

Claim 4.2.2 *Suppose that a path MACA x satisfies the unique deviator property. Then, there exists a (behavioral) Nash equilibrium y, such that $y(h) = x(h)$ for all $h \in H(x)$.*

To realize that the unique deviator property is necessary for the validity of this claim, recall Example 2.2 in Section 2. In the game depicted in Figure 2.2, (R_1, R_2) is a path MACA since it is supported by $Y \equiv \{R_1\} \times \{R_2\} \times \{L_3, R_3\}$. This path does not have the unique deviator property because both players 1 and 2 can deviate from it and reach player 3's information set. The path MACA (R_1, R_2) is not supported by (the unique) Nash equilibrium in this game.

A path MACA is also related to some of the solution concepts in the literature on "learning" (e.g., Fudenberg and Levine, 1993; Fudenberg and Kreps, 1995; Kalai and Lehrer, 1993). These notions are motivated by the fact that in repeated interactions among players, as players play equilibrium strategies, "off equilibrium choices" are not observed, and hence the requirement of commonality of beliefs about the actions that would have been taken in contingencies that were not realized during the play cannot be justified. Clearly, what players observe are actions taken along the equilibrium path of the play, and thus observed choices constitute a path CA. It turns out that viewed in this manner, a path MACA refines the notion of self-confirming equilibrium (Fudenberg and Levine, 1993) because an MACA requires that players' conjectures about their opponents be rational also "off the equilibrium path" (see Definition 3.2). Recall that in a self-confirming equilibrium (SCE), players learn only the path of play and their beliefs at information sets that are not reached can be arbitrary. Thus, a path MACA rules out many unreasonable outcomes. Dekel, Fudenberg, and Levine (1999) recently offered a refinement of SCE, which they call a "rationalizable self-confirming equilibrium" (RSCE). As the following claim asserts, our notion of a path MACA refines both concepts.

Claim 4.2.3 *Let x be a path MACA supported by the set Y. Then, every $y \in Y$ is a rationalizable self-confirming equilibrium (and, hence, a self-confirming equilibrium).*

We now provide two examples that show that path MACA strictly refines the set of rationalizable self-confirming equilibria. The reason is that we impose more stringent rationality restrictions on the players' reasonings. The first example, taken from Dekel, Fudenberg, and Levine (1999) is shown in Figure 4.2.4. As they correctly claim, outcome d_1 can arise from an RSCE in this game. However, d_1 is not a path MACA because this game has a unique subgame perfect equilibrium, (r_1, r_2, r_3), and thus this path is the unique path MACA. Example 4.1.7 is another demonstration of the fact that a path MACA provides a strict refinement of RSCE. Indeed, as we argued above, the unique path MACA in this game is (l_1, r_2), the path (r_1, r_2) is supported by an RSCE.

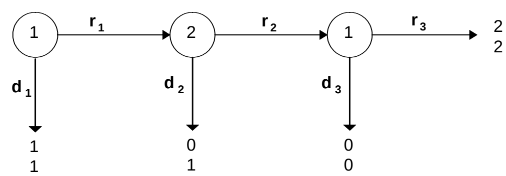

Figure 4.2.4

4.3 Null MACA

The null CA, i.e., the CA x where $x(h) = \emptyset$ for every information set $h \in H$, is the "coarsest" type of CA. This CA is suitable for analyzing situations where players do not communicate and have no common background (based on, say, past observations or social norms) that would imply an (implicit) agreement concerning the actions to be taken in some information sets. That is, players have no pre-conception about the course of action that might/should be taken, other than the one they can, individually and separately, deduce from the game. It is not surprising, therefore, that the null MACA turns out to be closely related to the notion of "rationalizable strategy profiles", due to Bernheim (1984) and Pearce (1984). Recall that in a normal form game, a rationalizable strategy for a player is one that is a best response to some beliefs he might have about his opponents' strategies, and beliefs are determined from the common knowledge of rationality (i.e., a player expects his opponents to also play only those strategies that are best responses to some beliefs they might have, and so on; see Appendix 1 for a formal definition). The null MACA therefore yields a rather intuitive refinement of (normal-form) rationalizability; in particular, an MACA, by Claim 3.3, excludes weakly dominated strategies.

Claim 4.3.1 *Suppose Y supports a null MACA. Then, every strategy profile in Y is (normal-form) rationalizable.*

For general extensive form games, there is no single accepted definition of rationalizability. The notion of a null MACA offers a new definition for this case, which, as the following example demonstrates, is quite attractive.

Example 4.3.2: Consider the game in Figure 4.3.2. In this game, the null MACA is supported by the set $Y = \{(l_1, (l_2, r_2))\}$. Hence, the only outcome of this game, in the absence of any contract, is the path l_1. This outcome is appealing because the strategy l_1 weakly dominates r_1. (Hence, by Claim 3.3, r_1 cannot belong to any set that supports an MACA.) In contrast, the well-known notions of normal-form and extensive-form rationalizability have no predictive power in this game. In particular, they support outcomes (r_1, l_2) and (r_1, r_2) as well.[15]

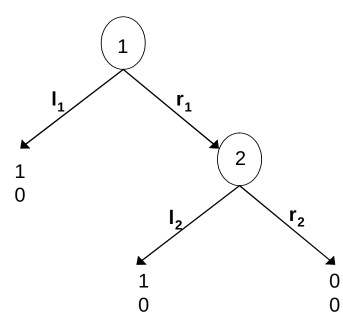

Figure 4.3.2

[15]There are other solution concepts that yield, in this example, the same outcomes as the null MACA; e.g., Asheim and Perea's (2003) "quasi-perfect rationalizability" and Dekel and Fudenberg's (1990) "permissibility".

5 Worlds Apart

We have seen so far how our framework helps analyze situations with incompletely specified courses of action that rational individuals may choose to follow for their own diverse reasons when the game being analyzed is common knowledge. We show in this section that our analysis can be easily extended to the more general case where players do not perceive the same game. (Like the proverbial set of blind men touching an elephant, each sees it differently, and believes that the way he sees it is the way the elephant is. Or like witnesses to an accident, each tells a complete story, but tells a different one. And, what other interpretation is there for "men are from Mars and women are from Venus"?)

The requirement that the "structure of the game" or the "model under discussion" be common knowledge seems to have been an integral part of game and economic theory. Yet, recent developments have started to doubt whether this is too strong an assumption, and whether a more plausible analysis could do with less. For example, Aumann and Brandenburger (1995; p. 1176-77) state the following regarding earlier assertions in the literature that common knowledge of the game or the model being analyzed is required for any economic analysis: "... This seemed sound when written, but in light of recent developments... it no longer does..... There is nothing about the real world that must be commonly known among the players."

As discussed in the Introduction, the situation where players live in different worlds cannot be described (hence, analyzed) by a single game that is common knowledge. Keeping within traditional noncooperative game theory, we formalize a single player's perception of a situation as Harsanyi's (1967-68) "types model" game, i.e., an (extensive form) game ${}^iT$.[16] Typically, the game ${}^iT$ is a game with imperfect information. This game incorporates all of i's beliefs about the other players. We do not assume that player i is certain about the "types" of the players he is facing. Rather, he may well be uncertain about some of the parameters (including the strategy spaces and the payoff functions of the other players) of the game he is (or thinks he is)

[16] We use ${}^iT$ rather than T^i in order to stress the fact that ${}^iT$ is not a restriction or projection of some underlying game T.

playing. Thus, ${}^iT$ completely describes player i's world.

Society, then, consists of n players, given by the set $N = \{1, 2, \ldots, n\}$ and n games $\{{}^1T, {}^2T, \ldots, {}^nT\}$. It is important to note that these n games may be very different; they need not even have the same structure. The classical paradigm where all players face a single game T that is common knowledge is obtained, in our framework, as a special case by imposing the (often unrealistic) restriction that for all $i, j \in N$, ${}^iT = {}^jT = T$. We analyzed this special case in the previous sections.

The generality of the notion of a course of action and of an MACA introduced in Definitions 3.1 and 3.2 allows us to straightforwardly extend the analysis to societies with diverse perceptions. The only modification that has to be introduced is that choices prescribed by a course of action must be "unambiguously interpretable" in all the n games.

Definition 5.1 *x is a course of action (CA) in society $\{{}^1T, {}^2T, \ldots, {}^nT\}$ if it is a CA in every game ${}^iT$, $i \in N$. A course of action x is mutually acceptable (MACA) in $\{{}^1T, {}^2T, \ldots, {}^nT\}$ if x is an MACA in every ${}^iT$, $i \in N$.*

While Definition 5.1 is a straightforward extension of Definition 3.2, note that it allows for a wide range of divergent views of the world to be modeled and analyzed. In particular, there is no limitation on the structure of different games ${}^iT$: these games could be very different depending on the extent to which norms, beliefs, and values are shared among different players. All that is required for shared action is that rational individuals follow that course of action for their own reasons (because it "makes sense" to them in their own world).

5.1 Applications in Different Worlds

5.1.1 Minmax Value

As stated in the Introduction, the notion that players may not analyze the same game is implicit in von Neumann and Morgenstern's (1947) justification of the minmax value in a two-person zero-sum game, as well as in Aumann and Brandenburger's (1995) defense of Nash equilibrium. We formally relate the minmax theorem to our framework below.

Let G be a two-person zero-sum game; Δ^i and $U^i, i = 1, 2$, are player i's (mixed) strategy sets and payoffs functions, respectively. Consider von Neumann and Morgenstern's argument, that player i assumes that player $j \neq i$ will somehow acquire the knowledge as to what player i's strategy is (see quote in Introduction). This implies in our framework that each player i analyzes the perfect information extensive form game ${}^iT$, where player i moves at the root of the game, followed by player $j \neq i$ who observes i's choice. Thus, players i and j see different games. For such games, the following result shows that every possible anticipation of player i, $i = 1, 2$, as to how the game might evolve results in i employing a Minmax strategy. Hence the outcome will be a Minmax play.

Claim 5.1.1 *Let Y support an MACA in ${}^iT$, $i = 1, 2$. Then, $y^i \in Y^i$ is player i's Minmax strategy in the game G.*[17]

The converse of Claim 5.1.1 does not hold. The reason is, again, that perfection imposes stricter conditions than Nash equilibrium. To realize this point, consider the two person zero-sum game in Figure 5.1.2(a), which is obtained by slightly modifying the game in Example 4.1.7. Observe that both (U, L) and (D, L) are Minmax (or Nash) equilibria. However, player 1's strategy U is weakly dominated by D.

	L	R
U	0,0	0,0
D	0,0	1,-1

Figure 5.1.2(a)

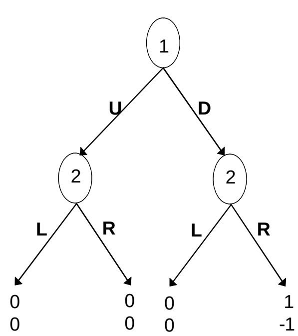

Figure 5.1.2(b)

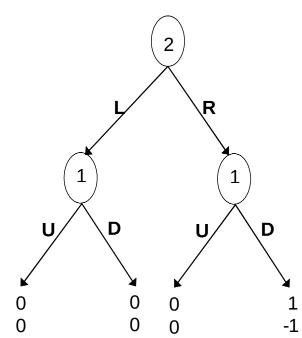

Figure 5.1.2(c)

[17] Observe that the strategy set of player i in ${}^iT$ is Δ^i. That is, in the game ${}^iT$, we have that $\mathbb{Y}^i = \Delta^i$.

Figures 5.1.2(b) and 5.1.2(c) depict respectively the games ${}^1T$ and ${}^2T$ associated with the normal form game in Figure 5.1.2(a). In game ${}^1T$, player 1 moves first, followed by player 2's move. The only choice by player 1 that can be supported in an MACA in game ${}^1T$ is D. Indeed, player 1's only best response if player 2 trembles is the strategy D, thus if Y supports an MACA in ${}^1T$, then every $y \in Y$ yields the path (D, L).

5.1.2 Military and Political Games

Examples abound in military conflicts and international politics where players perceived different games, and such situations readily lend themselves to formal analysis within our framework. During the second world war, the "Fall of France" to German army in 1940 is widely seen as a result of the Allied forces' inability to correctly perceive the likely German attack through the Ardennes. As discussed in several accounts of the war,[18] the Allied forces saw two possibilities for German attack on France: along the eastern border of France (the Maginot line), or an attack in the north through Belgium. In particular, the possibility of an attack through the Ardennes was rejected by the Allied forces as the Ardennes terrain was widely considered by the French to be unsuitable for tanks. German plans initially shared the French perceptions, and saw a long, costly encounter and a possible stalemate if the attack were carried out through Belgium or the Maginot line, as the Germans expected Allied forces to vigorously defend those two fronts. Being aware that Allied plans had rejected Ardennes as a possibility, Germans assigned a large force to attack through Ardennes and caught the Allied forces by surprise, leading to a quick victory for the German army.

Four years later, the Allied invasion of Normandy succeeded, and led to the end of the second World War, largely due to the successful attempts by the Allied forces to convince the Germans that the attack was to occur at Calais (see recent game-theoretic analyses of the Normandy invasion by Hendricks and McAfee, 2003, and Crawford, 2003) . In the Allied view, Calais was used as a 'feint' (complete with a mythical army group of 50 divisions, fake camps,

[18] See, e.g., Bennett and Dando (1979).

plywood airplanes and inflatable tanks) to mislead the German army as to the true site of the invasion. Germans believed Normandy was a feint, and that the true invasion was to come at Calais; so strong was this belief that the Germans directed all back-up forces to Calais and did not realize the main attack was occurring at Normandy until several days after the D-Day. Hendricks and McAfee (2003) analyze this situation, and the general incentive to feint, in a two-person signaling model. Crawford (2003) extends their model to include the possibility that players may be boundedly rational by assuming, in the context of a two-person matching pennies game, that players' types ('sophisticated' or 'mortal') are drawn from independent common knowledge distributions, and the game structure as well as payoffs are common knowledge.

Our framework offers a more general approach to study such situations. The "Fall of France" can be easily analyzed as different games viewed by France and Germany. This applies as well to the Normandy invasion, and similar economic situations where an 'attacker' has an incentive to feint and mislead the 'defender'. Another example, analyzed below, is the 1973 Middle East war.

Example 5.1.3: On October 6, 1973, Egypt attacked Israel, to Israel's complete surprise.[19] Since the 1967 war, Israel's military and almost all top Israeli officials had strongly held the belief that Egypt lacked the requisite military and air strength to launch an offensive against Israel, thus any movements of the Egyptian troops were viewed as 'routine manoeuvres' rather than preparations for war. With this belief, Israel chose not to mobilize its troops, since mobilization was costly. The Egyptian plans, on the other hand, sought to disguise its troop movements as manoeuvers, and set a target date of October 1^{st} to see whether Israel mobilized its troops in response. If Israel mobilized, the plan was to convert troop movements to normal manoeuvers. The games viewed by Israel and Egypt can be represented as in Figures 5.1.3(a) and (b), respectively. Note that these games have different structures, reflecting the fact that Israel was unaware of the October 1^{st} deadline. (The first payoff in the games below is for Egypt, second for Israel.)

[19] Our example follows Said and Hartley's (1982) description, who presented a game-theoretic view of the war.

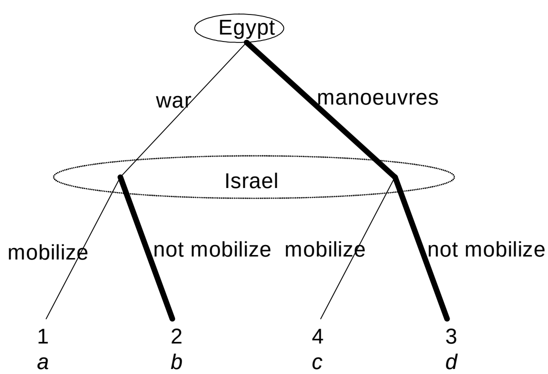

Figure 5.1.3(a): Israel's Game
(a>b, d>c)

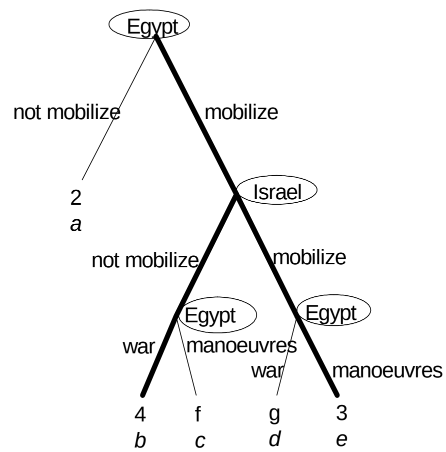

Figure 5.1.3(b): Egypt's Game
(a<c, f<4, g<2)

In the game viewed by Israel ($^{\text{Israel}}T$), it is a dominant strategy for Egypt not to go to war, hence, by Claim 3.3, the null MACA in this game predicts (*manoeuvers*, *not mobilize*). In Egypt's game ($^{\text{Egypt}}T$), however, Egypt would attack unless Israel mobilizes before October 1^{st}. Any strategy involving the choice (*not mobilize)* in Egypt's first information set is strictly dominated, so, by Claim 3.3, cannot be part of any MACA. Therefore, looking at the two games, we would expect Israel to not mobilize, and Egypt to mobilize and then go to war, as indeed was the case.

6 Conclusion

In this paper, we offered a framework to analyze situations where rational individuals with divergent perceptions agree to follow a common course of action. A course of action need not specify players' actions in all possible contingencies, and can be interpreted as a social norm or an incomplete contract. Players agree to a course of action if it makes sense to them in their own perceptions of the situation. Following traditional non-cooperative game theory, we first analyzed situations where all players perceive the same game, which is common knowledge. We showed that many different equilibrium concepts employed in economic analysis can be derived from our solution concept, a mutually acceptable course of action, by varying the degree to which the course of action specifies actions in various contingencies. In this setting, incomplete contracts emerge as rational responses by individuals with different beliefs and views of the world (or on the 'mistakes' other players may make), even without considering any transaction costs. Thus, our framework also serves to unify "game-theoretic" and "incomplete contracts" viewpoints. We then showed that the framework can easily be extended to situations where the players may have very different perceptions of the game they think they are playing. Our analysis provides one useful way of studying such situations which, we believe, are common in reality, as was demonstrated by the few examples of military conflicts and political games. Many other economic and social phenomena could be studied within our framework.

A Appendix 1: Definitions

Definition A1.1 *A perfect equilibrium is a strategy profile y with the property that there exists $y_k \rightsquigarrow y$ such that, for every player i, for all $h \in H^i$, for $k = 1, 2, \ldots$,*

$$u^i(y(h), y_k(-h)) \geq u^i(a, y_k(-h)) \text{ for all } a \in \Delta(h).$$

Definition A1.2 *A perfect Bayesian equilibrium is a strategy profile y with the property that for each player i and each information set $h \in H^i$, there exists a distribution $d(h)$ over the vertices in h such that: (i) the restriction of y^i to information sets that succeed h is i's best response to the restriction of y^{-i} to these information sets, using $d(h)$, and (ii) for any reachable information set h when y is followed, $d(h)$ is derived from y using Bayes' rule.*

Definition A1.3 *A strategy profile y is a rationalizable self-confirming equilibrium (RSCE) if there exists a collection of subsets of strategies $\{Y^i\}_{i \in N}$ such that, for every player i and every $z^i \in Y^i$, there exists a strategy z^j in the "extensive-form convex hull"*[20] *of Y^j (where $j \neq i$) such that*

1. *z has the same distribution over terminal nodes induced by y, and*
2. *for all information sets $h \in H^i$ that are not precluded by z^i, the restriction of z^i to information sets that succeed h is i's best response to the restriction of z^{-i} to these information sets, using the "sequential consistent" (w.r.t. z) distribution $d(h)$ over the vertices in h.*

Definition A1.4 *A collection of subsets of strategies $\{Y^i\}_{i \in N}$ is (normal-form) rationalizable if for every player i and every $y^i \in Y^i$, y^i is a best response to some profile in*

[20] A behavioral strategy y^j is in the extensive-form convex hull of Y^j, if there is an integer m, strategies $\{y^j_t\}_{t \in \{1,\ldots,m\}}$ in Y^j, sequences $y^j_{t,k} \rightsquigarrow y^j_t$, and a sequence $\lambda_k \rightarrow \lambda$ of probability distributions on $[1, ..., m]$, such that the strategies y^j_k generated by the convex combination of $y^j_{1,k}, ..., y^j_{m,k}$ with weights $\lambda_{1,k}, ..., \lambda_{m,k}$ converge to y^j (cf. Dekel et al. 2002, p. 476).

$\times_{j\neq i}\Delta\left(Y^j\right)$, where $\Delta\left(Y^j\right)$ denotes the set of all strategies $y^j \in Y^j$ such that y^j is outcome-equivalent[21] *to a distribution over Y^j.*

B Appendix 2: Proofs

Claim 0 *For any non-empty $Y^j \subseteq \mathbb{Y}^j$, $\Delta\left(Y^j\right) = \Delta^0\left(Y^j\right)$, where $\Delta\left(Y^j\right)$ [$\Delta^0\left(Y^j\right)$] denotes the set of all strategies $y^j \in \mathbb{Y}^j$ such that y^j is outcome-equivalent to a [finite-support] distribution over Y^j.*

Proof: Clearly, $\Delta\left(Y^j\right) \supseteq \Delta^0\left(Y^j\right)$. We need to show that $\Delta\left(Y^j\right) \subseteq \Delta^0\left(Y^j\right)$. Let $y^j \in \Delta\left(Y^j\right)$. Then, y^j is outcome-equivalent to a distribution μ over Y^j. Let ψ^j be the mixed strategy derived from the same distribution μ over $\varphi\left(Y^j\right) \equiv \left\{\varphi\left(y^j\right) \mid y^j \in Y^j\right\}$, where $\varphi\left(y^j\right)$ is the mixed representation of y^j.[22] Since $\varphi : Y^j \to \varphi\left(Y^j\right)$ is 1-1, by Kuhn's (1953) theorem and Pearce's (1984) Lemma 2, it is easy to see that ψ^j is outcome-equivalent to y^j. By using Pearce's (1984) Lemma 1, there exist $\varphi\left(y_t^j\right) \in \varphi\left(Y^j\right)$, $\lambda_t \geq 0$ $(t = 1, 2, \ldots, m)$, and $\sum_{t=1}^m \lambda_t = 1$ such that $\psi^j = \sum_{t=1}^m \lambda_t \varphi\left(y_t^j\right)$. Again by Kuhn's (1953) theorem, y^j is outcome-equivalent to the finite-support distribution λ over Y^j. *Q.E.D.*

Claim 3.3 *Suppose Y supports an MACA in a game T. Then, for each player i, $y^i \in Y^i$ is not a strictly dominated strategy. Moreover, if T is a game where each player has only one information set (in particular, if T represents a normal form game), then $y^i \in Y^i$ is not a weakly dominated strategy.*

Proof: For $y^i \in Y^i$ let $y_k^i \rightsquigarrow y^i$ and $y_k^j \overset{Y^j}{\rightsquigarrow} y^j$ (where $j \neq i$) satisfy the conditions in Definition 3.2. By condition (ii) in Definition 3.2 and continuity of u^i,

$$u^i(y(h), y(-h)) \geq u^i(a, y(-h)) \text{ for all } a \in \Delta(h).$$

[21]A strategy $y^j \in \mathbb{Y}^j$ is *outcome-equivalent* to a strategy $z^j \in \mathbb{Y}^j$ if for any $y^{-j} \in \mathbb{Y}^{-j}$ the strategy profiles (y^j, y^{-j}) and (z^j, y^{-j}) yield the same distribution over the terminal nodes.

[22]Recall that the mixed representation $\varphi\left(y^j\right)$ of a strategy y^j is the mixed strategy which assigns to a pure strategy s^j the product of all the probabilities assigned by y^j to the pure actions that define s^j.

By the one deviation property,[23] y^i is i's best response to y^{-i}. Thus, y^i is not a strictly dominated strategy.

Now, consider a game T where every player has only one information set. By condition (ii) in Definition 3.2, y^i is i's best response to y_k^{-i}. As for all $j \neq i$, y_k^j is a totally mixed strategy, it follows that y^i is not a weakly dominated strategy. *Q.E.D.*

Claim 3.4 *Let x be an MACA, and let y be a CA such that $y(h) = x(h)$ whenever $y(h) \neq \emptyset$. Then, y is an MACA.*

Proof: Since x is an MACA, there exists a set Y that supports x. Then, this Y also supports y. *Q.E.D.*

Claim 4.1.1 *The CA x is a complete MACA if and only if it is supported by the set $Y \equiv \{x\}$.*

Proof: Let Y support the complete MACA, x. For $y^i \in Y^i$ let $y_k^i \rightsquigarrow y^i$ and $y_k^j \overset{Y^j}{\rightsquigarrow} y^j$ (where $j \neq i$) satisfy the conditions in Definition 3.2. Since x is complete, condition (i) implies that $y_k^i \rightsquigarrow x^i$. Therefore, for every player i and every $y^i \in Y^i$, $y^i = x^i$. Hence, $Y = \{x\}$. *Q.E.D.*

Claim 4.1.2 *A strategy profile x is a complete MACA if and only if for each player i there exists a sequence $y_k[i] \rightsquigarrow x$ such that, for all $h \in H^i$ and for $k = 1, 2, \ldots,$*

$$u^i(x(h), y_k[i](-h)) \geq u^i(a, y_k[i](-h)) \text{ for all } a \in \Delta(h).$$

Proof: By Claim 4.1.1 we have that x is a complete MACA if and only if it is supported by the set $Y = \{x\}$. By Definition 3.2, therefore, x is a complete MACA if and only if for each player

[23] Let y be a behavioral strategy profile. According to the one deviation property (cf. Osborne and Rubinstein 1994, pp. 98-99), y_i is a best response to y_{-i} if the player i has no information set h at which a local change in $y(h)$ (holding the remainder of y fixed) increases his expected payoff.

i there exists a "trembling" sequence $\{y_k[i]\}_{k=1}^{\infty}$ such that (by condition (i)) $y_k[i] \rightsquigarrow x$, and (by condition (ii)), for all $h \in H^i$ and for $k = 1, 2, \ldots,$

$$u^i(x(h), y_k[i](-h)) \geq u^i(a, y_k[i](-h)) \text{ } for \text{ } all \text{ } a \in \Delta(h). \qquad Q.E.D.$$

Claim 4.1.3 *Every perfect equilibrium is a complete MACA.*

Proof: Let y be a perfect equilibrium. Then (see Definition A1.1) $Y = \{y\}$ supports y. Thus, y is a complete MACA. *Q.E.D.*

Claim 4.1.5 *A totally mixed complete CA is an MACA if and only if it is a perfect equilibrium.*

Proof: Let x be a totally mixed complete MACA. By Claim 4.1.1, it suffices to show x is a perfect equilibrium. By Claim 4.1.2 and the continuity of u^i we have that for all players i and for all $h \in H^i$,

$$u^i(x(h), x(-h)) \geq u^i(a, x(-h)) \text{ for all } a \in \Delta(h).$$

Since x is a totally mixed strategy, the sequence $y_k \equiv x$, $k = 1, 2, ..,$ can be chosen as the trembling sequence for every player i. Hence, x is a perfect equilibrium. *Q.E.D.*

Claim 4.1.6 *Every complete MACA is a perfect Bayesian equilibrium.*

Proof: Let x be a complete MACA. By Claim 4.1.2, for each player i, there exists a sequence $y_k[i] \rightsquigarrow x$. Define the distribution, $d^i(h)$, over the vertices in an information set $h \in H^i$ as follows: $d^i(h) \equiv \lim_{k\to\infty} d^i_k(h)$, where $d^i_k(h)$ is the unique distribution over these vertices derived from $y_k[i]$ using Bayes' rule. By condition (ii) in Definition 3.2 and the one deviation property, it therefore follows that player i's strategy x^i is a sequential best response to x^{-i} when i's beliefs over the vertices in $h \in H^i$ are given by $d^i(h)$. Hence, x is a perfect Bayesian equilibrium. *Q.E.D.*

Claim 4.2.2 *Suppose that a path MACA x satisfies the unique deviator property. Then, there exists a (behavioral) Nash equilibrium y, such that $y(h) = x(h)$ for all $h \in H(x)$.*

Proof: For each player i, let $y_k[i] \rightsquigarrow y[i]$ satisfy the conditions in Definition 3.2. Define the behavioral strategy y as follows:

$$y(h) \equiv \begin{cases} x(h), & \text{if } h \in H(x) \\ y[i(h)](h), & \text{if } h \notin H(x) \end{cases}.$$

By the unique deviator property, y is well defined. Condition (ii) of Definition 3.2 implies that y is a Nash equilibrium in behavioral strategies. *Q.E.D.*

Claim 4.2.3 *Let x be a path MACA supported by the set Y. Then, every $y \in Y$ is a rationalizable self-confirming equilibrium (and, hence, a self-confirming equilibrium).*

Proof: Let Y support a path MACA, x, and let $y^i \in Y^i$. Then, there exists a sequence $y_k[i] \rightsquigarrow y$ that satisfies the conditions in Definition 3.2. By condition (i) in Definition 3.2, the path generated by y coincides with x. Moreover, player i's strategy y^i is a sequential best response to y^{-i} when i's beliefs over the vertices in $h \in H^i$ are given by $d^i(h) \equiv \lim_{k\to\infty} d^i_k(h)$, where $d^i_k(h)$ is the unique distribution over these vertices derived from $y_k[i]$ using Bayes' rule (see proof of Claim 4.1.6). Therefore, for all information sets $h \in H^i$ that are not precluded by y^i, the restriction of y^i to information sets that succeed h is i's best response to the restriction of y^{-i} to these information sets, using the "sequential consistent" distribution $d(h)$ over the vertices in h.

To conclude the proof we have to show that for all $j \neq i$, y^j lies in the "extensive-form convex hull" of Y^j as defined by Dekel, Fudenberg and Levine (2002) (see Appendix 1). Indeed, for all $j \neq i$, $y^j_k[i] \overset{Y^j}{\rightsquigarrow} y^j$. That is, there exist $\{y^j_t\}_{t\in\{1,\ldots,m\}}$ in Y^j and $y^j_{t,k} \rightsquigarrow y^j_t$ $(t = 1, 2, \ldots, m)$ such that the behavioral strategies y^j_k generated by the convex combination of $y^j_{1,k}, \cdots, y^j_{m,k}$ with (the constant) weights $\lambda_1, \ldots.., \lambda_m$ converge to y^j. Hence y is a rationalizable self-confirming equilibrium. *Q.E.D.*

Claim 4.3.1 *Suppose Y supports a null MACA. Then, every strategy profile in Y is (normal-form) rationalizable.*

Proof: Fix a player i, for any $y^i \in Y^i$ let $y^i_k \rightsquigarrow y^i$ and $y^j_k \overset{Y^j}{\rightsquigarrow} y^j$ (where $j \neq i$) satisfy conditions in Definition 3.2. By condition (ii) of Definition 3.2 and the continuity of u^i, $u^i(y\,(h)\,, y\,(-h)) \geq u^i(a, y\,(-h))$ for all $h \in H^i$ and for all $a \in \Delta\,(h)$. By the one deviation property and Claim 0, y^i is a best response to some profile in $\times_{j \neq i} \Delta\left(Y^j\right)$. *Q.E.D.*

Claim 5.1.1 *Let Y support the null MACA in ${}^iT$, $i = 1, 2$. Then, $y^i \in Y^i$ is player i's Minmax strategy in the game G.*

Proof: Consider the game ${}^1T$. For $y^2 \in Y^2$ let $y_k \rightsquigarrow y$ satisfy the conditions in Definition 3.2. By condition (ii) of Definition 3.2, we have that $y(h)$ is player 2's best (mixed) action at every $h \in H^2$. Because the game is zero-sum, we have that for any $y^1 \in \mathbb{Y}^1$, $u^1(y^1, y^2) = Min_{b \in \Delta^2} U^1(y^1, b)$. It therefore follows that for every $y^2 \in \Delta(Y^2)$ and for any $y^1 \in \mathbb{Y}^1$, $u^1(y^1, y^2) = Min_{b \in \Delta^2} U^1(y^1, b)$.

Consider now $y^1 \in Y^1$ and let $y_k \rightsquigarrow y$ satisfy the conditions in Definition 3.2. By condition (ii) of Definition 3.2, we have that at the root of ${}^1T$ player 1's choice of the (mixed) action, y^1, is player 1's best response to y^2. That is, y^1 satisfies: $u^1(y^1, y^2) = Max_{y^1 \in \mathbb{Y}^1} u^1(y^1, y^2)$. Since $y^2 \in \Delta(Y^2)$, by our derivation above, we have that $u^1(y^1, y^2) = Max_{y^1 \in \mathbb{Y}^1} Min_{b \in \Delta^2} U^1(y^1, b) = Max_{a \in \Delta^1} Min_{b \in \Delta^2} U^1(a, b)$, i.e., y^1 is a Minmax strategy in G.

An analogous argument establishes that in the game ${}^2T$, y^2 is a Minmax strategy in G. *Q.E.D.*

References

[1] Asheim, G., A. Perea. 2003. Sequential and Quasi-Perfect Rationalizability in Extensive Games. mimeo.

[2] Aumann, R., A. Brandenburger. 1995. Epistemic Conditions for Nash Equilibrium. *Econometrica* **63**, 5, 1161-1180.

[3] Bennett, P. G., M. Dando. 1979. Complex Strategic Analysis: A Hypergame Study of the Fall of France. *Journal of the Operational Research Society* **30**, 23-32.

[4] Bernheim, D. 1984. Rationalizable Strategic Behavior. *Econometrica* **52**, 1007-1028.

[5] Bernheim, D. and M. Whinston. 1998. Incomplete Contracts and Strategic Ambiguity. *American Economic Review* **88**, 4, 902-932.

[6] Boudon, R. 1981. *The Logic of Social Action*. Routledge.

[7] Crawford, V. P. 2003. Lying for Strategic Advantage: Rational and Boundedly Rational Misrepresentation of Intentions. *American Economic Review*, **93**, 1, 133-149.

[8] Dekel, E., D. Fudenberg. 1990. Rational Behavior with Payoff Uncertainty. *Journal of Economic Theory* **52**, 243-67.

[9] Dekel, E., D. Fudenberg, D. Levine. 1999. Payoff Information and Self-Confirming Equilibrium. *Journal of Economic Theory* **89**, 165-185.

[10] Dekel, E., D. Fudenberg, D. Levine. 2002. Subjective Uncertainty over Behavior Strategies: A Correction. *Journal of Economic Theory* **104**, 473-478.

[11] Fudenberg, D., D. Kreps. 1995. Learning in Extensive Games I: Self-Confirming Equilibria. *Games and Economic Behavior* **8**, 20-55.

[12] Fudenberg, D., D. Levine. 1993. Self-Confirming Equilibrium. *Econometrica* **61**, 523- 546.

[13] Fudenberg, D., J. Tirole. 1991. Perfect Bayesian and Sequential Equilibrium. *Journal of Economic Theory* **53**, 236-260.

[14] Greenberg, J. 1990. *The Theory of Social Situations: An Alternative Game-Theoretic Approach.* Cambridge University Press.

[15] Greenberg, J. 2000. The Right to Remain Silent, *Theory and Decision* **48**, 193-204.

[16] Greenberg, J., D. Monderer, B. Shitovitz. 1996. Multistage Situations, *Econometrica* **64**, 1415-1437.

[17] Grossman, S., O. Hart. 1986. The Costs and Benefits of Ownership: A Theory of Vertical and Lateral Integration. *Journal of Political Economy* **94**, 691-719.

[18] Harsanyi, J. 1967-68. Games of Incomplete Information Played by Bayesian Players, Parts I-III. *Management Science* **14**, 159-182, 320-334, 486-502.

[19] Hart, O., J. Moore. 1988. Incomplete Contracts and Renegotiation. *Econometrica* **56**, 755-785.

[20] Hendricks, K., R.P. McAfee. 2003. Feints. mimeo.

[21] Kalai, E., E. Lehrer. 1993. Rational Learning Leads to Nash Equilibrium. *Econometrica* **61**, 1019-1045.

[22] Kreps, D. 1990. *Game Theory and Economic Modelling.* Oxford University Press.

[23] Kuhn, H. W. 1953. Extensive games and the problem of information, pp. 193-216 in: *Contributions to the Theory of Games, Volume II* (H. W. Kuhn and A. W. Tucker, eds.), Princeton University Press.

[24] Osborne, M. J., A. Rubinstein. 1994. *A Course in Game Theory.* The MIT Press, MA.

[25] Pearce, D. 1984. Rationalizable Strategic Behavior and the Problem of Perfection. *Econometrica* **52**, 1029-1050.

[26] Rubinstein, A. 1991. Comments on the Interpretation of Game Theory. *Econometrica* **59**, 909-924.

[27] Said, A. K., D. A. Hartley. 1982. A Hypergame Approach to Crisis Decision-Making: The 1973 Middle East War. *The Journal of the Operational Research Society* 33, 10, 937-948.

[28] Selten, R. 1975. Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games. *International Journal of Game Theory* **4**, 25-55.

[29] Snyder, L. L. 1979. *World in the Twentieth Century.* Krieger Publishing.

[30] von Neumann, J., O. Morgenstern. 1944. *Theory of Games and Economic Behavior.* Princeton University Press.

[31] Williamson, O. 1985. *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting.* The Free Press.