Beyond Normal Form Invariance:
First Mover Advantage in Two-Stage Games
with or without Predictable Cheap Talk

Peter J. Hammond

No 835

WARWICK ECONOMIC RESEARCH PAPERS

DEPARTMENT OF ECONOMICS

THE UNIVERSITY OF
WARWICK

Beyond Normal Form Invariance:
First Mover Advantage in Two-Stage Games
with or without Predictable Cheap Talk

Peter J. Hammond[*]

Final version: 10 January 2008

Abstract

Von Neumann (1928) not only introduced a fairly general version of
the extensive form game concept. He also hypothesized that only the
normal form was relevant to rational play. Yet even in Battle of the
Sexes, this hypothesis seems contradicted by players' actual behaviour
in experiments. Here a refined Nash equilibrium is proposed for games
where one player moves first, and the only other player moves second
without knowing the first move. The refinement relies on a tacit un-
derstanding of the only credible and straightforward perfect Bayesian
equilibrium in a corresponding game allowing a predictable direct form
of cheap talk.

[*]Department of Economics, University of Warwick, Coventry CV4 7AL, UK. e-mail:
`p.j.hammond@warwick.ac.uk`

1 Motivation and Introduction

1.1 Von Neumann's Standard Paradigm

Following Zermelo's (1912) pioneering analysis of chess and similar games, von Neumann (1928) devised a *standard paradigm* according to which multi-person decision problems in modern economic analysis and other social science are nearly always modelled as non-cooperative games in strategic form. This paradigm relies on two key assumptions, of which the first can be stated as follows:

Assumption 1. *A multi-person decision problem is fully described by a game in extensive form, whose structure is commonly known to all players in the game.*

Von Neumann's (1928) own extensive form description was later incorporated in *The Theory of Games and Economic Behavior*. Kuhn (1953) pointed out the implicit assumption that the order of different players' information sets was commonly known to all players at all stages of the game, and extended the von Neumann description to relax this assumption. Much more generally, we can now envisage an extensive form game as a stochastic process subject to the control of different players, with each player's information at each time described by a filtration. One key assumption, however, is that this stochastic process fits within Kolmogorov's (1933) framework, where everything random can be fitted within one overall probability space. As I have argued in Hammond (2007), this fails to allow for the possibility of having events that no player can foresee, and which may indeed even be impossible for any ideal observer to foresee.

1.2 Normal Form Invariance

The second assumption, which seems to have originated in von Neumann (1928), can be stated as follows:

Assumption 2. *It loses no generality to reduce the game in extensive form to the corresponding game in strategic or normal form, where each player makes a single strategic plan that covers all eventualities in the extensive form.*

It is perhaps worth going back all the way to von Neumann's original article, as adapted in von Neumann and Morgenstern (1943), in order to see how he justified normalizing the extensive form. First, normal form strategies are described on p. 79:

"Imagine now that each player . . . , instead of making each decision as the necessity for it arises, makes up his mind in advance for all possible contingencies; i.e., that the player . . . begins to play with a complete plan: a plan which specifies what choices he will make in every possible situation, for every possible actual information which he may possess at that moment in conformity with the pattern of information which the rules of the game provide for him for that case. We call such a plan a *strategy*."

Then pages 79–84 proceed to simplify the description of an extensive form game, to arrive at the normal form of the game in which each player makes just one move, and all moves are chosen simultaneously. In fact (p. 84):

"Each player must make his choice [of strategy] in absolute ignorance of the choices of the others. After all choices have been made, they are submitted to an umpire who determines . . . the outcome of the play for [each] player.
Observe that in this scheme no space is left for any kind of further 'strategy.' Each player has one move, and one move only; and he must make it in absolute ignorance of everything else."

Normalizing an extensive form game in this way is an extremely powerful device. And if the players of a game really do simultaneously submit their choices of a strategy to an umpire, who then sees that the players never deviate from their announced choices, then von Neumann and Morgenstern's claim on p. 85 seems entirely justified:

" . . . we obtained an all-inclusive formal characterization of the general game of n persons . . . . We followed up by developing an exact concept of strategy which permitted us to replace the rather complicated general scheme of a game by a much more simple special one, which was nevertheless shown to be fully equivalent to the former . . . . In the discussion which follows it will sometimes be more convenient to use one form, sometimes the other. It is therefore desirable to give them specific technical names. We will accordingly call them the *extensive* and the *normalized* form of the game, respectively.
Since these two forms are strictly equivalent, it is entirely within our province to use in each particular case whichever is technically more convenient at that moment. We propose, indeed, to make full use of this possibility, and must therefore re-emphasize that this does not in the least affect the absolute general validity of all our considerations."

It is this simplification that gives such power to familiar "normal form" concepts like Nash equilibrium, as well as to less familiar ones like trembling-hand perfect equilibrium (Selten, 1975), proper equilibrium (Myerson, 1978), correlated equilibrium (Aumann, 1987), rationalizable strategies (Bernheim, 1984; Pearce, 1984). Also, Mailath, Samuelson, and Swinkels (1993) show how even ostensibly extensive form ideas such as Selten's (1965) concept of subgame perfection, or Kreps and Wilson's (1982) concept of sequential equilibrium, have their (reduced) normal form counterparts.

Game theorists do relax normal form invariance somewhat by using extensive form solution concepts. For example, requiring players to respond credibly when other players deviate from expected behaviour was the original motivation for subgame perfection. See also Amershi, Sadanand and Sadanand (1985, 1989a, b, 1992), Hammond (1993), Sadanand and Sadanand (1995), Battigalli (1997), Battigalli and Siniscalchi (1999, 2002), and Asheim and Dufwenberg (2003), among other works that cast doubt on the normal form invariance hypothesis.

1.3   Outline of Paper

The purpose of this paper is to present a theoretical argument supporting the view that normal form invariance may be unduly restrictive. To do so, Section 2 considers a simple "Battle of the Sexes" game, where experimental evidence suggests that the first move does confer an advantage. It sets out the claim that this may be due to what would happen in the unique credible equilibrium of an associated game where cheap talk is possible after the first move, but before the second.

Section 3 begins to analyse a general two-stage game where one player moves first, and the only other player moves second, but without knowing the first player's move. It then allows simultaneous cheap talk by both players at an intermediate stage, between their two moves.

Because we are looking for an equilibrium that the players can infer, we require player 1's cheap talk to be "predictable" in the sense that it results from a pure strategy which is independent of her (hidden) action. Hence we consider a game where player 1 combines a mixed act with a pure message strategy. Afterwards player 2 first sends a message without knowing what 1 has done, then forms his conditional beliefs given 1's message and chooses an optimal mixed act accordingly.

Not surprisingly, any perfect Bayesian equilibrium (PBE) in the game with predictable cheap talk must induce a Nash equilibrium in the corresponding game without cheap talk. On the other hand, any Nash equilib-

rium without cheap talk can be extended into a PBE by making the second player "inattentive" to all cheap talk when forming his beliefs and choosing his strategy. Thus, cheap talk alone fails to refine the set of PBEs.

In order to facilitate such a refinement, Section 4 invokes a particular version of the revelation principle in the form due to Myerson (1982), as amended by Kumar (1985). First, this will allow player 2's message to be ignored, since anything he says could affect only his own beliefs and actions. Second, the revelation principle will allow general predictable cheap talk by player 1 to be replaced by "direct" cheap talk in the form of two suggestions for player 2, at his only ensuing information set: (i) the conditional probabilities that should be attached to player 1's earlier moves; (ii) player 2's choice of mixed act. Moreover, as argued in Section 4, we can limit attention to "straightforward" PBEs, where player 2 accepts both player 1's suggestions.

Section 5 finally introduces a credibility refinement. This requires a straightforward PBE to survive even when the player 2 is "Nash attentive" — i.e., when he accepts any suggestion by player 1 for choosing a Nash equilibrium of the game without cheap talk. The resulting "credible" equilibrium with cheap talk leads to an optimal Nash equilibrium for player 1 in the original game without cheap talk. When this optimal Nash equilibrium is unique, "sophistication" allows this cheap talk to remain implicit, so unnecessary. While these results may be hardly surprising, they do show how tacit communication can explain first-mover advantage in games like Battle of the Sexes.

Section 6 considers "virtual observability". This occurs when, as in Battle of the Sexes, sophistication effectively converts the game into one of perfect information, with the second player knowing the first move. Three examples show that virtual observability is rather special.

The concluding Section 7 discusses possible extensions and suggestions for future work that relaxes normal form invariance in other ways.

Except where it is standard, most notation will be explained wherever it is first used. Given any finite set F, however, let $\Delta(F)$ denote the set of probability distributions over F. Also, if F' is a proper subset of F, let $\Delta(F') \subset \Delta(F)$ denote those distributions that attach probability one to F'. Finally, if X and Y are arbitrary sets, let $X^Y := \prod_{y \in Y} X_y$ denote the set of all mappings $y \mapsto x_y$ from Y to X.

2 Battle of the Sexes

2.1 Two Different Extensive Forms

The two games in Figs. 1 and 2 are different extensive form versions of the familiar "Battle of the Sexes" game, whose normal form is given in Fig. 3. As is well known, there are two Nash equilibria in pure strategies, namely (B, b) and (S, s). There is also one mixed strategy Nash equilibrium where player 1 chooses B with probability $\frac{2}{3}$, and player 2 chooses b with probability $\frac{1}{3}$.

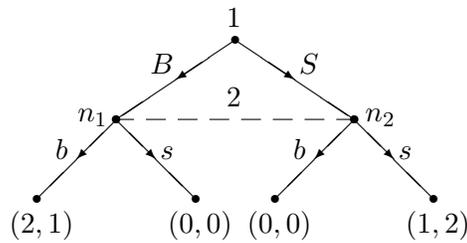Figure 1: Battle of the Sexes where Player 1 Moves First
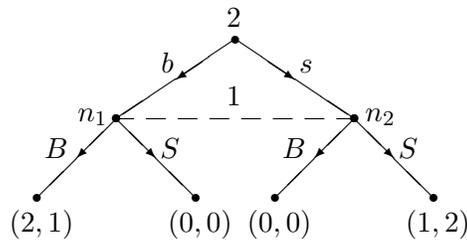
Figure 2: Battle of the Sexes where Player 2 Moves First

	b	s
B	2, 1	0, 0
S	0, 0	1, 2

Figure 3: Battle of the Sexes in Normal Form

Nevertheless, experiments strongly suggest that the player who moves first enjoys an advantage, in so far as (B, b) is played more often than (S, s)

in Fig. 1, but less often in Fig. 2.[1] These results have usually been ascribed to "positional order" or "presentation" effects that are seen as psychological or behavioral rather than fully rational responses to a change in the extensive form of the game.

2.2 Direct Cheap Talk in Battle of the Sexes

Consider the extensive form of Fig. 1, where player 1 moves first, and this is common knowledge. Suppose that, during an intermediate stage that succeeds player 1's move but precedes player 2's, the two players are allowed to communicate and indulge in unrestricted and mutually comprehensible "cheap talk".

As argued in Section 4, however, an extended version of the revelation principle implies that, in perfect Bayesian equilibrium (PBE) only player 1's cheap talk is relevant; it is already too late for player 2 to influence any action choice except his own. Moreover, we need only consider *direct* cheap talk where player 1's message m is a pair suggesting conditional probabilities $\rho(\cdot) \in \Delta(\{B, S\})$ and a mixed strategy $\sigma(\cdot) \in \Delta(\{b, s\})$ for player 2 at his only information set. Finally, the same principle allows us to limit attention to a "straightforward" PBE, where player 2 accepts 1's suggestions.

Now, any straightforward PBE would seem to involve just one of three possible direct messages that player 1 might send, corresponding to the three different Nash equilibria of the normal form:

1. corresponding to the equilibrium (B, b), a message with $\rho(B) = \sigma(b) = 1$ that yields the two players' expected payoffs of $(2, 1)$;

2. corresponding to the equilibrium (S, s), a message with $\rho(S) = \sigma(s) = 1$ that yields the two players' expected payoffs of $(1, 2)$;

3. corresponding to the mixed strategy equilibrium, a message with

$$\rho(B) = \sigma(s) = \tfrac{2}{3} \quad \text{and} \quad \rho(S) = \sigma(b) = \tfrac{1}{3}$$

---

[1]A "preliminary" experiment along these lines is described by Amershi, Sadanand, and Sadanand (1989b). Kreps (1990, p. 100) writes of "casual experiences playing this game with students". Later formal experiments yielding similar results were reported in Cooper *et al.* (1989, 1993). See also Schotter, Weigelt and Wilson (1994), Güth, Huck and Rapoport (1998), Muller and Sadanand (2003), Weber, Camerer and Knez (2004). The paper by Güth *et al.* even includes an experiment in which a form of cheap talk is explicitly allowed. The experimental design, however, includes the wording "B learns about A's decision" in the instructions. This may bias the results by offering the subjects too little encouragement to recognize the possibility of sending or receiving a deceptive message.

which yields the two players' expected payoffs of $(\frac{2}{3}, \frac{2}{3})$.

2.3 One Credible Equilibrium with Cheap Talk

In this Battle of the Sexes game with cheap talk, suppose all three "straightforward" messages could be regarded as credible. Then player 1 would expect player 2 to respond appropriately to whichever straightforward message she sends. So she would definitely choose the first of the three. But then, if player 2 hears any direct message except "I have played B and recommend that you play b", he should wonder whether player 1 has really not played B, or whether player 1 has somehow misspoken after playing B. Thus, player 2's best response to any other direct message actually becomes unclear. In the case of Battle of the Sexes, however, all that matters is that player 2 does choose b when player 1 suggests he should. This leaves us with just one possible outcome of any *credible* perfect Bayesian equilibrium (PBE).

Finally, if predictable direct cheap talk would produce a unique credible equilibrium message, we assume that both players are sufficiently "sophisticated" to reason what it will be. But this removes any need for cheap talk. Player 2 can work out the unique equilibrium message that he would receive in any credible PBE of the game with predictable direct cheap talk, and player 1 should know this also. By tacitly inferring what would happen if cheap talk were actually permitted, they reach the same unique outcome as in any credible PBE with predictable cheap talk.

3 General Two-Stage Games

3.1 The Basic Extensive Game

Instead of the specific Battle of the Sexes game discussed in Section 2, consider a general two-stage game Γ_0 with two players 1 and 2, for whom all the following facts are common knowledge. Player 1 begins the game by choosing an action a_1 from the finite set A_1. Then player 2 at his only information set, without seeing a_1, finishes the game by choosing an action a_2 from the finite set A_2. Each player i's payoff is denoted by $u_i(a_1, a_2)$ (for $i = 1, 2$). Allowing for mixed strategies $\alpha_i \in \Delta(A_i)$, the normal form of Γ_0 can be written as

$$G_0 = \langle \{1, 2\}, \Delta(A_1), \Delta(A_2), v_1, v_2 \rangle \tag{1}$$

with (expected) payoffs $v_i : \Delta(A_1) \times \Delta(A_2) \to \mathbb{R}$ for $i = 1, 2$ given by

$$v_i(\alpha_1, \alpha_2) := \sum_{a_1 \in A_1} \sum_{a_2 \in A_2} \alpha_1(a_1)\alpha_2(a_2)u_i(a_1, a_2). \tag{2}$$

Next, given their respective beliefs $\pi_1 \in \Delta(A_2)$ and $\pi_2 \in \Delta(A_1)$, define the two players' mixed strategy best response sets

$$B_1(\pi_1) \quad := \quad \underset{\alpha_1 \in \Delta(A_1)}{\arg\max} \, v_1(\alpha_1, \pi_1) \tag{3}$$

$$\text{and} \quad B_2(\pi_2) \quad := \quad \underset{\alpha_2 \in \Delta(A_2)}{\arg\max} \, v_2(\pi_2, \alpha_2). \tag{4}$$

Finally, we denote the set of mixed strategy Nash equilibra of G_0 by

$$E_0 := \{\, (\alpha_1, \alpha_2) \in \Delta(A_1) \times \Delta(A_2) \mid \alpha_1 \in B_1(\alpha_2), \ \alpha_2 \in B_2(\alpha_1) \,\}. \tag{5}$$

These are also the Nash (and perfect Bayesian) equilibra of Γ_0.

3.2 Predictable Cheap Talk

Cheap talk is introduced by allowing the two players to choose simultaneous message strategies $m_i \in M_i$ (for $i = 1, 2$) after player 1 has chosen a_1, but before player 2 chooses a_2. Often it will be convenient to let $m \in M := M_1 \times M_2$ denote the typical message pair (m_1, m_2). Of course, the main claim of this paper is precisely that it really *is* restrictive to reduce complex interactions to single strategy choices by each player.[2] Nevertheless, such restrictions seem not to detract from the force of the argument.

Also, we will look eventually for a predictable unique equilibrium of the game with cheap talk. Note, however, that no mixed message strategies could work this way; player 2 could not predict what messages result from such randomization. Indeed, player 1 cannot even make her message depend on the action that results from a mixed action strategy. So we consider only "predictable" cheap talk that results in one fixed message strategy for each player, independent of player 1's earlier action.

---

[2]Moreover, this rules out the kind of "long" cheap talk considered by Aumann and Hart (2003). Their model, however, involves messages that are sent by choosing one among only a finite set of "keystrokes". Also, the only example they provide of an equilibrium involving long cheap talk is presented in their Section 2.8. In a particular signalling game, it amounts to finding a mixed message strategy with infinite support. The formulation used here would allow any such message to be sent in only one stage.

3.3 An Extensive Form Game

An obvious two-person extensive game of perfect recall with predictable cheap talk proceeds in three successive stages as follows:

First action stage: Player 1 has one initial information set where she chooses a mixed action strategy $\alpha_1 \in \Delta(A_1)$.

Intermediate message stage: Both players simultaneously choose predictable messages $m_1 \in M_1$ and $m_2 \in M_2$. Though player 1 knows α_1 and even a_1, predictability rules out using this information. Hence both players communicate as though they have a single information set at this stage.

Second action stage: Player 2 has an information set $H_2(m)$ for each possible message pair $m \in M$. This enables him to choose a function $\alpha_2(\cdot|\cdot) \in [\Delta(A_2)]^M$ mapping each $m \in M$ to a mixed action strategy $\alpha_2(\cdot|m) \in \Delta(A_2)$.

Let Γ denote this extensive game. Its normal form can be written as

$$G = \langle \{1, 2\}, S_1, S_2, w_1, w_2 \rangle \tag{6}$$

where the two players' permitted (mixed) strategy sets have typical members denoted by

$$(\alpha_1, m_1) \in S_1 \quad := \quad \Delta(A_1) \times M_1 \tag{7}$$
$$\text{and} \quad (m_2, \alpha_2(\cdot|\cdot)) \in S_2 \quad := \quad M_2 \times [\Delta(A_2)]^M. \tag{8}$$

Also, definition (2) allows the two players' expected final payoffs $w_i : S_1 \times S_2 \to \mathbb{R}$ to be written as

$$w_i(\alpha_1, m_1, m_2, \alpha_2(\cdot|\cdot, \cdot)) := v_i(\alpha_1, \alpha_2(\cdot|m_1, m_2)). \tag{9}$$

3.4 Characterizing Perfect Bayesian Equilibrium

In a general extensive form game, a *perfect Bayesian equilibria* (PBE) is a *strategy–belief profile* which, for each player i and for each information set H where i has the move, combines: (i) a behavioural strategy specifying what (mixed) move i makes at H; (ii) a belief system specifying what subjective probabilities player i attaches to the different nodes of H. Moreover, this combination must satisfy the two requirements:

Consistent beliefs: Player i's beliefs at H are derived by Bayesian updating, provided the conditional probabilities are well defined given equilibrium moves at previous information sets;

Sequential rationality: Player i's move at H should maximize i's conditional expected payoff given the players' behavior strategies at all other information sets, and given player i's beliefs at H.

For the game Γ, accordingly, any strategy–belief profile involves player 2's conditional beliefs at each information set $H_2(m)$, after observing the message pair $m = (m_1, m_2) \in M$. We regard any such *belief system* as a mapping $m \mapsto \pi(\cdot|m)$ from M to $\Delta(A_1)$, denoted by

$$\pi(\cdot|\cdot) \in [\Delta(A_1)]^M. \tag{10}$$

We now give conditions for a particular strategy–belief profile

$$(\alpha_1^*, m^*, \alpha_2^*(\cdot|\cdot), \pi^*(\cdot|\cdot)) \in \Delta(A_1) \times M \times [\Delta(A_2)]^M \times [\Delta(A_1)]^M \tag{11}$$

in Γ to be a PBE.

At each last information set $H_2(m)$ of Γ, following the observed message pair $m \in M$, player 2's equilibrium belief system $\pi^*(\cdot|\cdot)$ determines his best response set $B_2(\pi^*(\cdot|m))$. Sequential rationality therefore requires player 2's behavior strategy at $H_2(m)$ to satisfy

$$\alpha_2^*(\cdot|m) \in B_2(\pi^*(\cdot|m)) \quad \text{for each } m \in M. \tag{12}$$

Earlier, anticipating player 2's equilibrium message m_2^* and sequentially rational response to each pair (m_1, m_2^*), player 1 chooses the pair

$$(\alpha_1^*, m_1^*) \in \underset{(\alpha_1, m_1) \in \Delta(A_1) \times M_1}{\arg\max} v_1(\alpha_1, \alpha_2^*(\cdot|m_1, m_2^*)). \tag{13}$$

This implies in particular that in the first action stage 1, anticipating both the equilibrium message pair $m^* \in M$ and player 2's induced response $\alpha_2^*(\cdot|m^*)$, player 1 chooses a mixed action strategy satisfying

$$\alpha_1^* \in B_1(\alpha_2^*(\cdot|m^*)). \tag{14}$$

During the intermediate message stage, player 2 anticipates player 1's choice of (α_1^*, m_1^*) and his own sequentially rational response to each pair $m \in M$. Hence player 2's equilibrium message m_2^* satisfies

$$m_2^* \in \underset{m_2 \in M_2}{\arg\max} v_2(\alpha_1^*, \alpha_2^*(\cdot|m_1^*, m_2)). \tag{15}$$

Finally, consistency of beliefs on the equilibrium path implies that

$$\pi^*(\cdot|m^*) = \alpha_1^*. \tag{16}$$

Then (12) implies that player 2 chooses a mixed strategy satisfying

$$\alpha_2^*(\cdot|m^*) \in B_2(\alpha_1^*). \tag{17}$$

3.5   Perfect Bayesian and Nash Equilibria

The following simple result establishes that, because any PBE of Γ induces Nash equilibrium strategies along an equilibrium path, it induces Nash equilibrium action strategies in the game G_0 without cheap talk.

Lemma 1. *Suppose the strategy–belief profile $(\alpha_1^*, m^*, \alpha_2^*(\cdot|\cdot), \pi^*(\cdot|\cdot))$ is a PBE in the game Γ with predictable cheap talk. Then the mixed action strategy profile $(\alpha_1^*, \alpha_2^*(\cdot|m^*))$ in $\Delta(A_1) \times \Delta(A_2)$ induced along the equilibrium path must be a Nash equilibrium in the game Γ_0 without cheap talk.*

Proof. Given the equilibrium message pair m^*, conditions (14) and (17) imply that the induced mixed strategies α_1^* and $\alpha_2^*(\cdot|m^*)$ are mutual best responses. So the strategy pair belongs to the set E_0 of Nash equilibria of the game Γ_0 without cheap talk, as defined in (5). $\square$

The next result shows that cheap talk alone excludes none of the Nash equilibria in the game Γ_0. In particular, all three Nash equilibria in the Battle of the Sexes example of Section 2 can be extended to PBEs with appropriate cheap talk.

Definition 1. *In the game Γ with predictable cheap talk, player 2's strategy–belief system $(\alpha_2(\cdot|\cdot), \pi(\cdot|\cdot)) \in [\Delta(A_2) \times \Delta(A_1)]^M$ is **inattentive** if both $\alpha_2(\cdot|m)$ and $\pi(\cdot|m)$ are constant, independent of m, for all message pairs $m \in M$. A PBE $(\alpha_1^*, m^*, \alpha_2^*(\cdot|\cdot), \pi^*(\cdot|\cdot))$ in Γ is **inattentive** if player 2's equilibrium strategy–belief system is inattentive.*

Lemma 2. *Let $(\bar{\alpha}_1, \bar{\alpha}_2) \in E_0$ be any Nash equilibrium in the game Γ_0 without cheap talk. Let M be any message space for player 1. Then the corresponding game Γ with predictable cheap talk in M has an inattentive PBE which induces $(\bar{\alpha}_1, \bar{\alpha}_2)$ along the equilibrium path.*

Proof. Consider the strategy–belief profile in Γ where:

1. player 1 combines $\alpha_1^* = \bar{\alpha}_1$ with an arbitrary message $m_1^* \in M_1$;

2. player 2 sends an arbitrary message $m_2^* \in M_2$;

3. player 2's strategy–belief system is inattentive, with

$$\alpha_2^*(\cdot|m) = \bar{\alpha}_2 \quad \text{and} \quad \pi^*(\cdot|m) = \bar{\alpha}_1 \quad \text{for all } m \in M. \qquad (18)$$

It is easy to see that $(\alpha_1^*, m^*, \alpha_2^*(\cdot|\cdot), \pi^*(\cdot|\cdot))$ must be a PBE. $\qquad \square$

4 An Extended Revelation Principle

4.1 Direct Cheap Talk

The revelation principle will involve a new game $\hat{\Gamma}$ which is like Γ except that:

1. Player 2's message space M_2 becomes a singleton $\{\bar{m}_2\}$, so he can only send a constant message $\bar{m}_2$. This makes 2's message irrelevant, of course, so we ignore it from now on.

2. Player 1's general messages $m_1 \in M_1$ are replaced by *direct* messages

$$\hat{m} = (\rho, \sigma) \in \hat{M} := \Delta(A_1) \times \Delta(A_2). \qquad (19)$$

Here, following Kumar's (1985) extension of the revelation principle, the first component $\rho \in \Delta(A_1)$ of each direct message that player 1 might send can be interpreted as beliefs about player 1's strategy that 1 suggests to 2. Following Myerson (1982), the second component $\sigma \in \Delta(A_2)$ can be interpreted as the mixed strategy that 1 suggests to 2.[3]

The typical strategy–belief profile in the game $\hat{\Gamma}$ with direct cheap talk will be denoted by

$$(\hat{\alpha}_1, \hat{m}, \hat{\alpha}_2(\cdot|\cdot), \hat{\pi}(\cdot|\cdot) \in \Delta(A_1) \times \hat{M} \times [\Delta(A_2)]^{\hat{M}} \times [\Delta(A_1)]^{\hat{M}}. \qquad (20)$$

---

[3]Following Forges (1986), many later writers describe direct messages as "canonical".

4.2   Equivalent Straightforward Equilibria

Definition 2. *In the game $\hat{\Gamma}$ with direct cheap talk, the strategy–belief profile $(\hat{\alpha}_1, \hat{m}, \hat{\alpha}_2(\cdot|\cdot), \hat{\pi}(\cdot|\cdot))$ with $\hat{m} = (\rho, \sigma)$ is **straightforward** if*

$$\hat{\pi}(\cdot|\hat{m}) = \rho = \hat{\alpha}_1 \quad and \quad \hat{\alpha}_2(\cdot|\hat{m}) = \sigma. \tag{21}$$

*A strategy–belief profile that is straightforward and also a PBE is a **straightforward PBE**.*

That is, a strategy–belief profile is straightforward if player 1 suggests beliefs that match her mixed action and if player 2 accepts both suggestions that make up player 1's direct message.

The following result extends to our setting the versions of the revelation principle due to Myerson (1982) and Kumar (1985).

Theorem 1. *Let $(\alpha_1^*, m^*, \alpha_2^*(\cdot|\cdot), \pi^*(\cdot|\cdot))$ be any PBE strategy–belief profile in the game Γ with general predictable cheap talk. Then in the associated game $\hat{\Gamma}$ with direct cheap talk there is an **equivalent** PBE*

$$(\hat{\alpha}_1^*, \hat{m}^*, \hat{\alpha}_2^*(\cdot|\cdot), \hat{\pi}^*(\cdot|\cdot)) \tag{22}$$

that is inattentive, straightforward, and generates the same equilibrium action strategy pair

$$(\hat{\alpha}_1^*, \hat{\alpha}_2^*(\cdot|\hat{m}^*)) = (\alpha_1^*, \alpha_2^*(\cdot|m^*)). \tag{23}$$

Proof. By Lemma 1, the mixed action strategy pair $(\alpha_1^*, \alpha_2^*(\cdot|m^*))$ generated by the PBE of Γ must be a Nash equilibrium of the game Γ_0 without cheap talk. To construct the equivalent PBE strategy–belief profile (22), first choose $\hat{\alpha}_1^* = \alpha_1^*$. Next, define the equivalent direct message $\hat{m}^* \in \hat{M}$ in the game $\hat{\Gamma}$ as the Nash equilibrium pair $(\alpha_1^*, \alpha_2^*(\cdot|m^*))$ itself. Finally, define an inattentive strategy–belief system for player 2 by choosing $\hat{\pi}^*(\cdot|\hat{m}) := \alpha_1^*$ and $\hat{\alpha}_2^*(\cdot|\hat{m}) := \alpha_2^*(\cdot|m^*)$ for each direct message $\hat{m} \in \hat{M} = \Delta(A_1) \times \Delta(A_2)$.

Evidently the constructed strategy–belief profile (22) is both inattentive and straightforward. As in Lemma 2, it is also a PBE of $\hat{\Gamma}$. $\square$

The extended revelation principle is especially useful in allowing any PBE in the game Γ with predictable cheap talk to be converted to an inattentive straightforward PBE in the associated game $\hat{\Gamma}$ with direct cheap talk. Nevertheless, Lemma 2 applies even in $\hat{\Gamma}$. For this reason, an extra consideration is needed to refine the set of Nash equilibria.

5 Credible Equilibria with Direct Cheap Talk

5.1 Nash Attentiveness

The following definition requires player 2 to accept player 1's direct message in $\hat{\Gamma}$ whenever it suggests a specific Nash equilibrium of the game Γ_0 without cheap talk.

Definition 3. *In the game $\hat{\Gamma}$ with direct cheap talk, player 2's strategy–belief system $(\hat{\alpha}_2(\cdot|\cdot), \hat{\pi}(\cdot|\cdot)) \in [\Delta(A_2) \times \Delta(A_1)]^{\hat{M}}$ is **Nash attentive** if it satisfies $(\hat{\alpha}_2(\cdot|\hat{m}), \hat{\pi}(\cdot|\hat{m})) = \hat{m}$ whenever the direct message $\hat{m} = (\rho, \sigma) \in \hat{M} = \Delta(A_1) \times \Delta(A_2)$, viewed as a pair of mixed strategies, constitutes a Nash equilibrium of the game Γ_0 without cheap talk. A PBE strategy–belief profile is **Nash attentive** if player 2's strategy–belief system is Nash attentive.*

5.2 First-Mover Advantage with Cheap Talk

We now show that the PBEs of $\hat{\Gamma}$ with Nash attentive beliefs generate Nash equilibria in Γ_0 that are optimal for the first mover.

Definition 4. *In the game Γ_0 without cheap talk, the Nash equilibrium mixed strategy pair $(\alpha_1^*, \alpha_2^*) \in \Delta(A_1) \times \Delta(A_2)$ is **optimal for player 1** if $v_1(\alpha_1^*, \alpha_2^*) \geq v_1(\alpha_1, \alpha_2)$ for all (α_1, α_2) in the set E_0 of Nash equilibria in Γ_0. The same pair is **uniquely optimal for player 1** if $v_1(\alpha_1^*, \alpha_2^*) > v_1(\alpha_1, \alpha_2)$ for all alternative Nash equilibria $(\alpha_1, \alpha_2) \in E_0 \setminus \{(\alpha_1^*, \alpha_2^*)\}$.*

Theorem 2. *Let $(\hat{\alpha}_1^*, \hat{m}^*, \hat{\alpha}_2^*(\cdot|\cdot), \hat{\pi}^*(\cdot|\cdot))$ be any straightforward Nash attentive PBE strategy–belief profile in the game $\hat{\Gamma}$ with predictable direct cheap talk. Then the action profile $(\alpha_1^*, \alpha_2^*) := (\hat{\alpha}_1^*, \hat{\alpha}_2^*(\cdot|\hat{m}^*))$ induced on the equilibrium path is an optimal Nash equilibrium for player 1 in the game Γ_0 without cheap talk.*

Proof. Applying equilibrium condition (13) to $\hat{\Gamma}$ instead of Γ gives

$$(\hat{\alpha}_1^*, \hat{m}^*) \in \underset{(\alpha_1, \hat{m}) \in \Delta(A_1) \times \hat{M}}{\arg\max} v_1(\alpha_1, \hat{\alpha}_2^*(\cdot|\hat{m})). \tag{24}$$

Let $(\bar{\alpha}_1, \bar{\alpha}_2) \in E_0$ be any Nash equilibrium in Γ_0. Because player 2's strategy $\hat{\alpha}_2^*(\cdot|\hat{m})$ is Nash attentive in the game $\hat{\Gamma}$, player 1's expected payoff from choosing $(\alpha_1, \hat{m})$ with $\alpha_1 = \bar{\alpha}_1$ and $\hat{m} = (\bar{\alpha}_1, \bar{\alpha}_2)$ will be

$$v_1(\bar{\alpha}_1, \hat{\alpha}_2^*(\cdot|\hat{m})) = v_1(\bar{\alpha}_1, \bar{\alpha}_2). \tag{25}$$

14

Now (24) implies that $v_1(\hat{\alpha}_1^*, \hat{\alpha}_2^*(\cdot|\hat{m}^*)) \geq v_1(\bar{\alpha}_1, \hat{\alpha}_2^*(\cdot|\hat{m}))$, and so $v_1(\alpha_1^*, \alpha_2^*) \geq v_1(\bar{\alpha}_1, \bar{\alpha}_2)$ by (25). This holds for every $(\bar{\alpha}_1, \bar{\alpha}_2) \in E_0$. But Lemma 1 implies that $(\alpha_1^*, \alpha_2^*) \in E_0$, so it must be an optimal Nash equilibrium for player 1. $\square$

The next definition considers what happens when player 2 may not be fully Nash attentive, but is nevertheless attentive at least to messages that suggest following a Nash attentive straightforward PBE.

Definition 5. *A straightforward PBE strategy–belief profile in the game $\hat{G}$ with direct cheap talk is* **credible** *if it is identical to a Nash attentive straightforward PBE along the equilibrium path.*

Obviously, by Theorem 2, any such credible PBE must also induce an optimal Nash equilibrium outcome for player 1.

5.3   First-Mover Advantage without Cheap Talk

Suppose the game $\hat{\Gamma}$ with predictable direct cheap talk has a unique credible PBE. Then the two players can reasonably expect each other to infer what this direct cheap talk would be, even in the game Γ_0 without cheap talk. The following definition singles out the corresponding Nash equilibrium of this game.

Definition 6. *A Nash equilibrium of the game Γ_0 without cheap talk is* **sophisticated** *if it is induced by a credible straightforward PBE of the corresponding game $\hat{\Gamma}$ with predictable direct cheap talk, and moreover this credible PBE is unique.*

Theorem 3. *Suppose (α_1^*, α_2^*) is a uniquely optimal Nash equilibrium for player 1 in Γ_0. Then (α_1^*, α_2^*) is the unique sophisticated equilibrium.*

Proof. Theorem 2 implies that there is a unique credible PBE of $\hat{\Gamma}$, and that this equilibrium induces (α_1^*, α_2^*). $\square$

Figure 4 specifies an example of a normal form game G_0 in which, if player 1 moves first in the associated extensive form Γ_0, there is no sophisticated equilibrium. Not surprisingly, cheap talk plays a key role here in enabling coordination on one of the two Nash equilibria that are equally good for player 1. But if the two players' payoffs after (L, ℓ) were $(1 + \epsilon, \delta)$ instead, for any $\epsilon > 0$ and any $\delta > 0$, then (L, ℓ) would be the unique sophisticated equilibrium.

	ℓ	r
L	$1,1$	$0,0$
R	$0,0$	$1,1$

Figure 4: A Game with No Sophisticated Equilibrium

6 The Special Case of Virtual Observability

6.1 Definition

Corresponding to our basic game Γ_0 without cheap talk, there is an associated extensive form game

$$\Gamma_1 := \langle \{1,2\}, \Delta(A_1), [\Delta(A_2)]^{A_1}, v_1, v_2 \rangle \qquad (26)$$

of perfect information, where player 2 is informed of 1's move and so can make his mixed strategy $\alpha_2 \in \Delta(A_2)$ a function of player 1's action a_1. Now the Battle of Sexes example of Fig. 1 has a unique sophisticated equilibrium where both players effectively act as though player 1's move could indeed be observed. It is a case where the same pure strategy profile $(a_1, a_2) \in A_1 \times A_2$ in the game G_0 happens to be both the unique outcome of any credible PBE in $\hat{\Gamma}$ and of any subgame perfect equilibrium in Γ_1. Weber, Camerer and Knez (2004) call this "virtual observability". The next three examples remind us that it is really a very special property.

6.2 Duopoly: Cournot versus Stackelberg

Consider a duopoly where firm 1 is able to choose its quantity before firm 2. Also, suppose both firms know this and that firm 2 can observe 1's output. Then it is fairly obvious that any sophisticated equilibrium must be a subgame perfect equilibrium where firm 1 acts as a Stackelberg leader and firm 2 as a follower. If firm 1's output remains hidden, however, the normal form of the game corresponds to one in which the duopolists choose their quantities simultaneously. Then a sophisticated equilibrium is Cournot.

For example, suppose each firm $i \in \{1, 2\}$ has the profit function

$$\Pi_i(q_i, q_j) = \beta_i q_i - \gamma q_i q_j - \tfrac{1}{2} q_i^2$$

which is quadratic in its own quantity q_i and also depends on the other's quantity q_j. Suppose too that each firm is risk neutral and so maximizes expected profit. Finally, suppose that the three parameters β_1, β_2 and γ

are positive and satisfy the restrictions $\beta_1 > \gamma\beta_2$, $\beta_2 > \gamma\beta_1$, and $\gamma < 1/\sqrt{2}$. Even if the first firm pursues a mixed strategy, the second firm's optimal choice satisfies $q_2 = \beta_2 - \gamma\,\mathbb{E}\,q_1$, where $\mathbb{E}$ denotes the mathematical expectation. Thus, the first firm's expected profit is

$$\mathbb{E}\Pi_1 = (\alpha_1 - \gamma q_2)\mathbb{E}q_1 + \gamma^2(\mathbb{E}q_1)^2 - \tfrac{1}{2}\mathbb{E}q_1^2.$$

This is maximized by choosing the Stackelberg leader's pure strategy $q_1^S := (\beta_1 - \gamma\beta_2)/(1 - 2\gamma^2)$, which exceeds the unique Cournot equilibrium quantity $q_1^C := (\beta_1 - \gamma\beta_2)/(1 - \gamma^2)$. It follows that virtual observability fails, even though there is a unique Nash equilibrium and it uses pure strategies.

6.3 Mixed Strategies

Consider the simple and familiar example of matching pennies, whose normal form is shown in Figure 5. There is a unique Nash equilibrium, associated with a unique straightforward PBE strategy–belief profile in the corresponding game of predictable direct cheap talk. The only direct message $\hat{m} = (\rho, \sigma) \in \Delta(\{H, T\}) \times \Delta(\{h, t\})$ that is sent in this unique equilibrium has $\rho(H) = \rho(T) = \sigma(h) = \sigma(t) = \tfrac{1}{2}$. Obviously, the need for mixed action strategies in Nash equilibrium implies that virtual observability cannot hold.

	h	t
H	$1, -1$	$-1, 1$
T	$-1, 1$	$1, -1$

Figure 5: Matching Pennies

6.4 Multiple Nash Equilibria

	h	t	e
H	$4, 0$	$0, 4$	$0, 0$
T	$0, 4$	$4, 0$	$0, 0$
E	$0, 0$	$0, 0$	$1, 1$

Figure 6: Extended Matching Pennies

The game in Fig. 6 is matching pennies played for a stake of \$4 supplied by a third party. The game is also extended by allowing each (steady handed)

player to choose "edge" as well as heads or tails. If just one player chooses edge, the stake is withdrawn, and neither wins anything. But if both choose edge the third party pays each \$1 for being imaginative.

In the corresponding extensive form game Γ_1 with perfect information where player 1 moves first, player 2 would choose: t in response to H; h in response to T; and e in response to E. So Γ_1 has (E, e) as a unique subgame perfect equilibrium. This is not induced by a credible straightforward PBE of $\hat{G}$, however, because a better Nash equilibrium of G_0 for player 1 is the familiar mixed strategy equilibrium with $\alpha_1(H) = \alpha_1(T) = \alpha_2(h) = \alpha_2(t) = \frac{1}{2}$, since player 1's expected payoff is 2 rather than 1. Once again, virtual observability fails, and in this case it does so even though the unique subgame perfect equilibrium is a Nash equilibrium in pure strategies.

6.5   Implications of Virtual Unobservability

When virtual observability fails, the extensive game Γ_0 is fundamentally different from Γ_1 where player 2 is informed of player 1's earlier move. Sometimes, as in Figs. 5 and 6, this is because player 1 gains by keeping her initial move concealed. Sometimes, however, as in Section 6.2, player 1 could gain from having her initial move revealed. In that example, the first duopolist would earn more profit from being a Stackelberg leader. It would also like to report having chosen the Stackelberg leader's optimal quantity q_1^S, expecting the second firm to choose its best response $q_2^S := \beta_2 - \gamma q_1^S$. However, that report is not credible because, if it were believed, the first firm does even better by choosing its best response $q_1 = \beta_1 - \gamma q_2^S \neq q_1^S$. So requiring the follower to be attentive only to the Nash equilibrium message q_1^C in any Nash attentive straightforward PBE imposes a binding constraint on the leader's strategy choice.

7   Concluding Remarks

7.1   Beyond Experimental Anomalies

Experimental economists have recognized that there is a first-mover advantage in Battle of the Sexes and similar games. They typically ascribe this advantage, however, to "positional" or "presentational" effects, suggesting the need to look beyond orthodox rationality concepts in order to explain their experimental results.

This paper, by contrast, introduces a "sophisticated" refinement of Nash equilibrium that can explain first-mover advantage using only a minor varia-

tion of standard rationality and equilibrium concepts. This refinement, like the "manipulated Nash equilibrium" concept explored in Amershi, Sadanand and Sadanand (1985, 1989a, b, 1992) and in Sadanand and Sadanand (1995), depends on the extensive form of the game. So it violates von Neumann's hypothesis of normal form invariance. Unlike manipulated Nash equilibrium, however, the tacit communication that underlies forward induction arguments is explicitly modelled through a corresponding game with cheap talk. This cheap talk is required to be predictable so that it can remain tacit.[4]

Nevertheless, the precise relationship between sophisticated and manipulated Nash equilibrium deserves further exploration. The ideas presented here should also be applied to a much broader class of games, starting with the "recursive games" considered in Hammond (1982).

7.2 Beyond Orthodox Game Theory

Much of orthodox game theory is built on two assumptions of what one may call the "ZNK paradigm" — due to Zermelo (1912), won Neumann (1928) and Kolmogorov (1933). This paper has criticized normal form invariance, the second of these. But the first, claiming that games can be modelled with a single extensive form, is also questionable, as discussed in Hammond (2007). So, of course, is a third key assumption, namely that all players are fully rational, and so will always find the optimal action at each information set.

Indeed, following Zermelo (1912), orthodox game theory predicts that any two-person zero-sum game of perfect information such as Go should be played perfectly, and so perfectly predictably. Yet we find the following in a prominent novel by an author who won the Nobel Prize for Literature in 1968.

> " 'This is what war must be like,' said Iwamoto gravely.
>
> He meant of course that in actual battle the unforeseeable occurs and fates are sealed in an instant. Such were the implications of White 130. All the plans and studies of the players, all the predictions of us amateurs and of the professionals as well had been sent flying.
>
> As an amateur, I did not immediately see that White 130 assured the defeat of the 'invincible Master.' "

---

[4]A conjecture is that relaxing predictability in the game with cheap talk would allow player 1 to achieve her optimal correlated equilibrium. Where this is better than her optimal Nash equilibrium, cheap talk is essential as a correlation device. Without it, player 2 cannot infer what correlated equilibrium strategy to choose.

Yasunari Kawabata (1954) *The Master of Go*, translated from the author's own shortened version by Edward G. Seidensticker (New York: Alfred A. Knopf, 1972); end of chapter 37.

Such considerations remind us how far the three standard assumptions take us from reality. To conclude, it seems that the systematic study of games and economic behaviour has barely progressed beyond a promising but possibly misleading beginning.

References

Amershi, A.H., Sadanand, A.B., and Sadanand, V. (1985) "Manipulated Nash Equilibria I: Forward Induction and Thought Process Dynamics in Extensive Form" University of British Columbia, Economics Working Paper 928.

Amershi, A.H., Sadanand, A.B, and Sadanand, V. (1989a) "Manipulated Nash Equilibria II: Some Properties" University of Guelph, Economics Discussion Paper 1989-5.

Amershi, A.H., Sadanand, A.B, and Sadanand, V. (1989b) "Manipulated Nash Equilibria III: Applications and a Preliminary Experiment" University of Guelph, Economics Discussion Paper 1989-6.

Amershi A.H., A. Sadanand and V. Sadanand (1992) "Player Importance and Forward Induction" *Economics Letters* 38, 291–297.

Asheim, G. and M. Dufwenberg (2003) "Deductive Reasoning in Extensive Games" *Economic Journal* 113, 305–325.

Aumann, R.J. (1987) "Correlated Equilibrium as an Expression of Bayesian Rationality" *Econometrica* 55, 1–18.

Aumann, R. J. and S. Hart (2003) "Long Cheap Talk" *Econometrica*, 71, 1619–1660.

Battigalli, P. (1997) "On Rationalizability in Extensive Games" *Journal of Economic Theory* 74, 40–60.

Battigalli, P. and D. Siniscalchi (1999) "Hierarchies of Conditional Beliefs and Interactive Epistemology in Dynamic games" *Journal of Economic Theory* 88, 188–230.

Battigalli, P. and D. Siniscalchi (2002) "Strong Belief and Forward-Induction Reasoning" *Journal of Economic Theory* 106, 356–391.

Bernheim, B.D. (1984) "Rationalizable Strategic Behavior" *Econometrica* 52, 1007–1028.

Cooper, R., Dejong, D.V., Forsythe, R., and Ross, T.W. (1989) "Communication in the Battle-of-the-Sexes Games: Some Experimental Evidence" *Rand Journal of Economics* 20, 568–587.

Cooper, R., Dejong, D.V., Forsythe, R., and Ross, T.W. (1993) "Forward Induction in the Battle-of-the-Sexes Games" *American Economic Review* 83, 1303–1315.

Forges, F. (1986) "An Approach to Communication Equilibria" *Econometrica* 54, 1375–1385.

Güth, W., Huck, S., and Rapoport, A. (1998) "The Limitations of the Positional Order Effect: Can it Support Silent Threats and Non-Equilibrium Behavior?" *Journal of Economic Behavior and Organization* 34, 313–325.

Hammond, P. J. (1982) "Sophisticated Dynamic Equilibria for Extensive Games" Economics Technical Report, Institute for Mathematical Studies in the Social Sciences, Stanford University.

Hammond, P. J. (1993) "Aspects of Rationalizable Behavior" in K. Binmore, A.P. Kirman, and P. Tani (eds.) *Frontiers of Game Theory* (Cambridge, Mass.: M.I.T Press, 1993), ch. 14, pp. 277–305.

Hammond, P. J. (2007) "Schumpeterian Innovation in Modelling Decisions, Games, and Economic Behaviour" *History of Economic Ideas* 15, 179–195.

Kolmogorov, A.N. (1933, 1956) *Grundbegriffe der Wahrscheinlichkeitsrechnung* (Berlin: Springer); translated as *Foundations of Probability* (New York: Chelsea).

Kreps, D. (1990) *Game Theory and Economic Modelling* (Oxford: Oxford University Press).

Kreps, D. M. and R. Wilson (1982) "Sequential Equilibrium" *Econometrica* 50, 863–894.

Kuhn, H.W. (1953) "Extensive Games and the Problem of Information" *Contributions to the Theory of Games, II*, 193–216; reprinted in H.W. Kuhn (ed.) *Classics in Game Theory* (Princeton: Princeton University Press, 1997).

Kumar, P. (1985) "Consistent Mechanism Design and the Noisy Revelation Principle" in *Essays on Intertemporal Incentives and Signalling* Ph.D. dissertation, Dept. of Economics, Stanford University.

Mailath, G.J., L. Samuelson and J.M. Swinkels (1993) "Extensive Form Reasoning in Normal Form Games" *Econometrica* 61, 273–302.

Muller, R.A. and A. Sadanand (2003) "Order of Play, Forward Induction, and Presentation Effects in Two-Person Games" *Experimental Economics* 6, 5–25.

Myerson, R.B. (1978) "Refinements of the Nash Equilibrium Concept" *International Journal of Game Theory* 7, 73–80.

Myerson, R. (1982) "Optimal Coordination Mechanisms in Generalized Principal–Agent Problems" *Journal of Mathematical Economics* 10, 67–81.

Pearce, D. (1984) "Rationalizable Strategic Behavior and the Problem of Perfection" *Econometrica* 52, 1029–1050.

Rapoport, A. (1997) "Order of Play in Strategically Equivalent Games in Extensive Form" *International Journal of Game Theory* 26, 113–136.

Sadanand, A. and V. Sadanand (1995) "Equilibria in Non-cooperative Games II: Deviations Based Refinements of Nash Equilibrium" *Bulletin of Economic Research* 47, 93–113.

Schotter, A., Weigelt, K., and Wilson, C. (1994) "A Laboratory Investigation of Multiperson Rationality and Presentation Effects" *Games and Economic Behavior* 6, 445–468.

Selten, R. (1965) "Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit" *Zeitschrift für die gesamte Staatswissenschaft* 121, 301–324 and 667–689.

Selten, R. (1975) "Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games" *International Journal of Game Theory* 4, 25–55.

Von Neumann, J. (1928) "Zur Theorie der Gesellschaftsspiele" *Mathematische Annalen*, 100, 295–320; reprinted in A.H. Taub (ed.) *Collected Works of John von Neumann, Vol. VI* (Oxford: Pergamon Press, 1963), pp. 1–26; translated as "On the Theory of Games of Strategy" in A.W. Tucker and R.D. Luce (eds.) *Contributions to the Theory of Games, Vol. IV* (Princeton: Princeton University Press, 1959), pp. 13–42.

Von Neumann, J. and O. Morgenstern (1943, 1953) *Theory of Games and Economic Behavior* 3rd. ed. (Princeton: Princeton University Press).

Weber, R.A., Camerer, C.F., and Knez, M. (2004) "Timing and Virtual Observability in Ultimatum Bargaining and 'Weak Link' Coordination Games" *Experimental Economics* 7: 25–48.

Zermelo, E. (1912) "Über eine Anwendung der Mengenlehre auf die Theorie des Schachspiels" *Proceedings of the Fifth International Congress of Mathematicians*, Vol. II, pp. 501–504.