Reinterpreting Mixed Strategy Equilibria: A Unification of the Classical and Bayesian Views*

Philip J. Reny
Department of Economics
University of Chicago

Arthur J. Robson
Department of Economics
University of Western Ontario

September 2002

Abstract

We provide a new interpretation of mixed strategy equilibria that incorporates both von Neumann and Morgenstern's classical concealment role of mixing as well as the more recent Bayesian view originating with Harsanyi. For any two-person game, G, we consider an incomplete information game, $\mathcal{IG}$, in which each player's type is the probability he assigns to the event that his mixed strategy in G is "found out" by his opponent. We show that, generically, any regular equilibrium of G can be approximated by an equilibrium of $\mathcal{IG}$ in which almost every type of each player is strictly optimizing. This leads us to interpret i's equilibrium mixed strategy in G as a combination of deliberate randomization by i together with uncertainty on j's part about which randomization i will employ. We also show that such randomization is not unusual: For example, i's randomization is nondegenerate whenever the support of an equilibrium contains cyclic best replies.

*We wish to thank Drew Fudenberg, Motty Perry and Hugo Sonnenschein for very helpful comments. We are especially grateful to Bob Aumann for a stimulating discussion of the literature on mixed strategies that prompted us to fit our work into that literature and ultimately led us to the interpretation we provide here, and to Hari Govindan whose detailed suggestions permitted a substantial simplification of the proof of our main result. Reny gratefully acknowledges financial support from the National Science Foundation (SES-9905599). Robson appreciates support from the Social Sciences and Humanities Research Council of Canada and from a Canada Council for the Arts Killam Research Fellowship.

1. Introduction

The purpose of this paper is to better understand mixed strategy Nash equilibria in finite two-person games. In particular, we show that a player's equilibrium mixture can be usefully understood partly in terms of deliberate randomization by the player, and partly as an expression of the opponent's uncertainty about which randomization the player will employ. This allows us to unify the otherwise sharply distinct views of mixed strategies proposed by von Neumann and Morgenstern (1944) for zero-sum games and by Harsanyi (1973) for nonzero-sum games.

Von Neumann and Morgenstern (vNM (1944)), when focusing on two-person zero-sum games, unequivocally interpret mixing as an act of deliberate randomization whose purpose is to conceal. They point out that each player strictly prefers any one of his equilibrium strategies over any other strategy if he is certain that the mixed strategy he chooses will be found out by his opponent prior to his opponent's choice.

For example, in Matching Pennies player 1 strictly prefers the fifty-fifty mixture over every other mixed strategy when he knows that player 2 will find out the mixed strategy he chooses. VNM conclude from this that there is a defensive or concealment rationale for mixing in zero-sum games:

Thus one important consideration for a player in such a game is to protect himself against having his intentions found out by his opponent. Playing several such strategies at random, so that only their probabilities are determined is a very effective way to achieve a degree of such protection: By this device the opponent cannot possibly find out what the player's strategy is going to be, since the player does not know it himself (vNM (1944, p. 146)).

According to the *classical* rationale, then, a mixed strategy represents deliberate randomization on a player's part. Everyone, including the player himself, is uncertain about that player's pure choice. However, because it is based on the desirability of concealment, the classical rationale for mixing runs into difficulties in nonzero-sum games. As Schelling notes:

The essence of randomization in a two-person zero-sum game is to preclude the adversary from gaining intelligence about one's mode of play... In games that mix conflict with common interest, however, randomization plays no such central role... (Schelling (1960, p. 175)).

Consider, for example, the mixed equilibrium in the Battle of the Sexes. In this equilibrium, *neither* player can be thought of as deliberately randomizing to conceal his pure choice because, in this game, *each* player prefers to reveal his pure choice, whatever it is, to the other player. Thus, the classical rationale is inappropriate here.

But if concealment is not the rationale for mixing in general games, what is? Harsanyi (1973) provides an ingenious answer. He shows that virtually any mixed equilibrium can be viewed as an equilibrium of a nearby game of incomplete information in which small private variations in the players' payoffs lead them to strictly prefer one of their pure strategies. Thus, from Harsanyi's point of view no player ever actually randomizes and a player i is uncertain about another player j's pure choice only because it varies with j's type, which is private information. The significant conceptual idea introduced by Harsanyi is that a player's mixed strategy expresses the ignorance of the *other* players—*not of the player himself*—about that player's pure strategy choice.[1]

Aumann (1987; see especially Section 6) takes Harsanyi's idea further. By eliminating the payoff perturbations altogether, he directly interprets a player's mixed strategy solely as an expression of the other players' uncertainty about that player's pure strategy choice. This view is now widespread.[2] Indeed, as Aumann and Brandenburger (1995) remark:

In recent years, a different view of mixing has emerged. According to this view, players do not randomize; each player chooses some definite action. But other players need not know which one, and the mixture represents their uncertainty, their conjecture about his choice.

Thus, the view of mixing that has emerged, the *Bayesian view* let us call it, bears no resemblance to the classical view that mixing represents a deliberate decision to randomize in order to conceal one's choice. The concealment role of mixing has been entirely left behind.

In contrast, we argue here that the intuitively appealing classical view can be incorporated into a general interpretation of mixed equilibria. In fact, the approach we propose is tied to *both* the Bayesian and classical views.

[1] We thank Bob Aumann for suggesting to us that this *conceptual* contribution by Haranyi was at least as important as his formal purification theorem.

[2] For example, see Armbruster and Boege (1979), Tan and Werlang (1988) and Brandenburger and Dekel (1989).

Our approach is tied to the Bayesian view by incorporating Harsanyi's idea that a player's private information can lead to uncertainty about that player's choice from the opponent's perspective. However, our approach differs crucially from Harsanyi's in terms of the precise nature of the players' private information. In our setup, there is *no uncertainty about payoffs*. Rather, each player is concerned that his opponent might find out his choice of mixed strategy, and it is *the level of this concern* that is private information. As in Harsanyi, such private information can make the opponent more uncertain about a player's choice than the player himself. But, in our approach, because each player is concerned that his mixed strategy might be found out by his opponent, he may benefit from the concealment effect of deliberate randomization. This simultaneously ties our approach to the classical view.

More precisely, we interpret equilibria of any finite two-person game G as limits of equilibria of certain sufficiently nearby games of incomplete information, $\mathcal{IG}$. Each incomplete information game, $\mathcal{IG}$, is derived from G as follows. Nature moves first by independently choosing, for each player i, a type, $t_i \in [0, 1]$, according to some continuous distribution. Each player is privately informed of his own type, which is his assessment of the probability that the opponent will find out his *mixed strategy* before the opponent moves. We shall be concerned with the equilibria of $\mathcal{IG}$ as the type distributions become concentrated around zero and so as the players' concerns for being found out vanish.[3,4]

As in Harsanyi, our model provides the players with strict incentives. Indeed, as we show, any regular equilibrium of G can be approximated by an equilibrium of $\mathcal{IG}$ in which almost every type of each player is strictly optimizing. But there is an important difference. Harsanyi's players strictly prefer to use only pure strategies, while our players in general strictly prefer to use *mixed* strategies. When our players mix, they do so deliberately, because the benefits of concealment make

[3] The reader may wonder why each player is concerned that the other might find out his mixed strategy when neither player considers it possible that he will find out the other's mixed strategy. The answer is that the player's decision problem, when he finds out the other's mixture, is trivial – he simply best replies. Hence, these decision problems are supressed. A formal resolution along these lines is provided in Appendix A.

[4] The possibility that one's opponent finds out one's mixed strategy arises naturally in a repeated game setting, where an opponent can deduce one's stage game mixed strategy by repeatedly observing the *pure* outcomes of one's randomized play. Such a repeated game interpretation of our model is pursued in Section 9, where we provide a dynamic game in which only *past pure actions* are observable and where the analysis of the dynamic game nonetheless can be reduced to an analysis of our static game $\mathcal{IG}$, in which *mixed* strategies are sometimes observed.

this strictly optimal, not because the equilibrium requires them to make the other player indifferent. Conversely, when our players choose pure strategies, they do so because randomization is harmful and they actively wish to reveal their choice to the other player.

For example, in the unique equilibrium of our incomplete information perturbation, $\mathcal{IG}$, of Matching Pennies, all types of both players strictly optimize by choosing the fifty-fifty mixture (see Section 2). Thus, neither player's behavior depends upon his private information and each player deliberately randomizes. Such randomization is strictly beneficial because each player believes the other might find out his mixed strategy. Our approach therefore supports the classical view of the mixed equilibrium in Matching Pennies, namely, that each player deliberately randomizes fifty-fifty and is certain that his opponent will do so as well. Indeed, Theorem 4.1 generalizes this to all zero-sum games.

On the other hand, all equilibria of $\mathcal{IG}$ near the completely mixed equilibrium of Battle of the Sexes require almost every type of each player to employ one of his two pure strategies (see Section 2). Concealment is shown to play no role in this equilibrium precisely because each player prefers to reveal his pure choice in this game. Moreover, our rationale for mixing here coincides with the Bayesian view: Each player i employs one or the other pure strategy; player j does not know which pure strategy i will employ, but assigns some probability to each one, where these probabilities are given by i's equilibrium mixture. This is generalized in Theorem 6.5 which states that if, starting from any cell in G's payoff matrix, both players' payoffs increase whenever either one of them switches to a best reply, then every equilibrium of our incomplete information perturbation requires almost every type of each player to employ a pure strategy.

So, our interpretation coincides with the classical view in zero-sum games, and it coincides with the Bayesian view in a class of coordination games. But what about the vast majority of games lying between these two extremes? As shown by example in Section 2, our interpretation will typically differ from both the Bayesian and classical views. The example is a 3x3 nonzero-sum game with a unique completely mixed equilibrium, m^*. In our incomplete information perturbation, no type of either player chooses his completely mixed equilibrium strategy, yet no type of either player chooses a pure strategy either. Instead, almost every type of each player i strictly optimizes by using one of three mixed strategies, m_{i1}, m_{i2}, or m_{i3}, each of which gives positive weight to just two pure strategies. Each randomization, m_{ik}, benefits i by optimally concealing the two pure strategies in its support. Further, if μ_{ik} denotes the fraction of player i's types using

m_{ik} in the limit as the players' concerns for being found out converge to zero, then $m_i^* = \mu_{i1}m_{i1} + \mu_{i2}m_{i2} + \mu_{i3}m_{i3}$.

The above three games serve to exemplify the following general interpretation of any equilibrium, m^*, of the original game G :

Each player i's equilibrium mixture, m_i^*, can be expressed as a convex combination of a fixed finite set of i's mixed strategies, m_{ik}, say. Each mixed strategy in the convex combination represents a strategy that i might deliberately employ, while the weight on that mixed strategy represents the opponent's belief that i will employ it

Such convex combinations reveal the role of deliberate randomization. In our perturbed game, where players are slightly concerned that their mixed strategy might be found out, the strategies m_{ik} are strictly optimal for the types employing them and, when the m_{ik} are non-degenerate, they optimally conceal the pure strategies in their support.

From this perspective, the classical and Bayesian views are extreme. On the one hand, our interpretation coincides with the classical view only when the above convex combination is degenerate, placing full weight on i's equilibrium mixed strategy, as in matching pennies. On the other, our interpretation coincides with the Bayesian view only when every mixed strategy in the above convex combination is pure, with weights given by i's equilibrium mixture, as in the Battle of the Sexes. In general, our interpretation differs from both the classical and Bayesian views, as typified by the third example above.

A strength of our interpretation is that it eliminates the sharp distinction between zero-sum and nonzero-sum games insofar as the role of randomization is concerned. For example, according to our view, when moving from matching pennies to the battle of the sexes through continuous payoff changes, the role played by deliberate randomization in their mixed equilibria continuously diminishes to zero.

One might wonder when at least one of the m_{ik} above is nondegenerate, because then our interpretation involves deliberate randomization. Theorem 6.2 states that if the support of an equilibrium, m^*, of any game G contains a best-reply cycle, then a positive (and bounded away from zero) measure of both players' types must use non-degenerate mixed strategies, m_{ik}, in any approximating equilibrium of $\mathcal{IG}$. Hence, the presence of best reply cycles in the support of an equilibrium of a two-person game indicates a role for deliberate randomization in that equilibrium.

The key decomposition of m_i^* into a convex combination of the m_{ik} can be obtained without considering the, often complex, limiting equilibria of $\mathcal{IG}$. Theorem 7.1 states that such decompositions are solutions to a linear optimization problem.

We restrict attention throughout to two-person games. Additional issues arise with more players. For example, with three or more players one must specify how many opponents find out a player's mixed strategy. There does not appear to be a single natural choice here. However, there is no reason to doubt that any reasonable choice will yield strict incentives to mix in some games, as we obtain here.

In addition to the work cited above, a rich literature on purification has grown out of Harsanyi's (1973) seminal contribution. (See, for example, Radner and Rosenthal (1982) and Aumann et al. (1983).) The central issue in this literature is whether every mixed strategy equilibrium of an incomplete information game is (perhaps approximately) equivalent to some pure strategy equilibrium. In our model, this is not an issue because, as we shall show, all equilibria of $\mathcal{IG}$ are pure, generically. But note that a pure strategy in $\mathcal{IG}$ allows the players' to choose non-degenerate mixed strategies from G.

More closely related are Rosenthal (1991) and Robson (1994).[5] Both papers are concerned with the robustness of equilibria of two-person games to changes in the information structure. Rosenthal observes that equilibria of some nonzero-sum games remain equilibria even when the opponent is sure to find out one's mixed strategy choice.

Robson perturbs arbitrary two-person games by supposing that each player's pure or mixed strategy is found out by the opponent with a common known probability. Equilibria that survive arbitrarily small perturbations of this kind are called "informationally robust." Robson shows that informationally robust equilibria exist and refine Nash equilibria.[6] He also observes that informational robustness with respect to mixed strategies yields strict incentives to mix in some 2x2 nonzero-sum examples. However, in a typical informationally robust equilibrium, the players will not have strict incentives and they randomize in order to make the opponent indifferent.

[5] Less closely related is Matsui (1989). He considers a repeated game with a small probability that one player's entire supergame strategy will be revealed to the other before this second player chooses. This implies that any subgame perfect equilibrium is Pareto-efficient.

[6] Our results here imply that, generically, the sets of informationally robust equilibria and Nash equilibria coincide.

The remainder of the paper is organized as follows. Section 2 contains three leading examples illustrating the main ideas. Section 3 describes our incomplete information perturbation of an arbitrary two-person game. Section 4 provides our results concerning zero-sum games, while Section 5 analyzes the more challenging nonzero-sum case and contains our main approximation theorem. Section 6 provides conditions under which our interpretation necessarily involves the classical view that players deliberately randomize, as well as a condition under which our interpretation involves only the Bayesian view in which no player randomizes. Section 7 provides the decomposition theorem outlined above, and Section 8 provides an example showing the potential significance of unused strategies. Finally, Section 9 shows how our static model, in which players find out their opponent's *mixed* strategy with some probability, is the reduced form of a dynamic game in which two randomly matched populations play a stage game repeatedly and where a small fraction of players observe their opponent's past history of *pure* actions.

2. Three Leading Examples

The scope of the present approach can be demonstrated by considering three examples: Matching Pennies, Battle of the Sexes, and Modified Rock-Scissors-Paper. To each of these normal form games, G, say, we associate a nearby game of incomplete information, $\mathcal{IG}$, which we now describe informally.

For $0 \leq \underline{\varepsilon} < \bar{\varepsilon} \leq 1$, the players' types, t_1 for player 1 and t_2 for player 2, are drawn independently and uniformly from $[\underline{\varepsilon}, \bar{\varepsilon}]$.[7] The players choose a mixed strategy in G as a function of their type.[8] With probability $1-t_i$, player i receives the payoff in G from the pair of mixed strategies chosen, whereas with probability t_i, he receives the payoff in G resulting from his mixed strategy choice together with a best reply against it.[9,10]

Accordingly, we interpret a player's type to be the probability he assigns to the event that the opponent finds out his mixed strategy and best replies to

[7] The general results here hold for a much broader class of distributions.

[8] These are pure strategies in $\mathcal{IG}$. Players are, as usual, also allowed to employ mixed strategies. But as we shall see, mixed strategies in $\mathcal{IG}$ are typically suboptimal.

[9] If there are multiple best replies for j against i's mixed strategy, then one that is best for i is employed. See Section 3.

[10] If player i uses a mixed strategy in $\mathcal{IG}$ specifying, for each of his types, a randomization device, or "lottery," for choosing a mixed strategy in G, then with probability t_i his payoff is determined by the mixed strategy that is the outcome of the lottery together with a best reply against it.

	H	T
H	1,-1	-1, 1
T	-1, 1	1,-1

Figure 2.1: Matching Pennies

it. However, note that neither player believes he will find out the opponent's mixed strategy.[11] Hence each type of each player makes only the single decision of choosing a mixed strategy in G. Note also that player i cares only about the overall distribution over pure strategies in G induced by the opponent's strategy in $\mathcal{IG}$. This is because, from i's point of view, the opponent's strategy in $\mathcal{IG}$ is relevant for determining i's payoff only when the opponent does not find out i's mixed strategy.[12]

We are interested in the limiting equilibria of $\mathcal{IG}$ as $\bar{\varepsilon}$ and $\underline{\varepsilon}$ tend to zero, so that $\mathcal{IG}$ tends to the original game G.

2.1. Matching Pennies

Recall vNM's observation that in Matching Pennies (Figure 2.1) the players strictly prefer the fifty-fifty mixture when they are sure to be found out. We shall show that for every $\bar{\varepsilon} > \underline{\varepsilon} > 0$, including those near zero, $\mathcal{IG}$ has an equilibrium in which every type of each player chooses to randomize fifty-fifty over H and T and that this non-degenerate mixture is strictly optimal.[13]

So, suppose that every type of one player, player 2, say, uses the fifty-fifty mixture. It suffices to show that the fifty-fifty mixture is the uniquely optimal choice for every type of the other player, player 1, say.

Figure 2.2 shows player 1's payoff as a function of the probability, p, that his mixed strategy assigns to H, given that player 2 finds out this mixed strategy.

[11] For a strategically equivalent description in which it is common knowledge that each player might find out the other's mixed strategy, see Appendix A.

[12] See Section 3.

[13] In fact, this is the essentially unique equilibrium, as follows from the general theorem of Section 4.

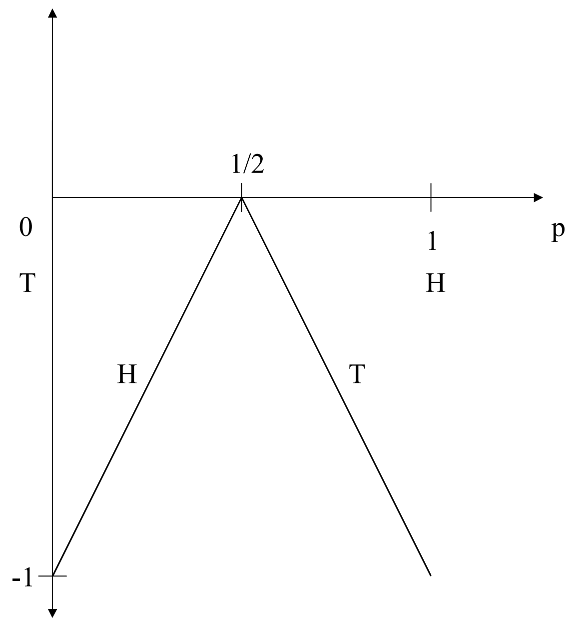

Figure 2.2: Player 1 Found Out in Matching Pennies

Player 1's payoff is negative both for $p \in [0, 1/2)$, where 2's best reply is H, and for $p \in (1/2, 1]$ where 2's best reply is T, and so is uniquely maximized at $p = 1/2$, where it is zero, regardless of 2's reply.

Now consider player 1 in $\mathcal{IG}$ when his type, the probability he assigns to being found out, is $t_1 \in [\underline{\varepsilon}, \bar{\varepsilon}]$. Because the other player is mixing equally, any type of player 1 is indifferent among all his mixed strategies *conditional on not being found out*. Hence, because player 1 of type $t_1 > 0$ assigns positive probability to the event that he *is* found out, the fifty-fifty mixture is the uniquely optimal choice for every type of player 1, as claimed.

Thus, the incomplete information game associated with Matching Pennies captures the classical point of view that mixing is a deliberate attempt to conceal one's choice. Consider now the Battle of the Sexes, whose mixed equilibrium cannot be rationalized by a concealment argument. We now show that its associated incomplete information game *can* rationalize the mixed equilibrium and, rather than conceal them, the players actively reveal their pure choices.

	L	R
T	2,1	0,0
B	0,0	1,2

Figure 2.3: Battle of the Sexes

2.2. Battle of the Sexes

Consider the Battle of the Sexes (henceforth BoS; see Figure 2.3) and the mixed equilibrium in which each player chooses his favorite pure strategy with probability 2/3. (His "favorite" pure strategy yields him a payoff of 2 if coordination is achieved.)

For this example, set $\underline{\varepsilon} = 0$, so that in $\mathcal{IG}$ the player types are drawn independently and uniformly from $[0, \bar{\varepsilon}]$. We will show that for $\bar{\varepsilon} > 0$ small enough, there is an equilibrium of $\mathcal{IG}$ in which approximately 2/3 of each player's types choose that player's favorite BoS pure strategy and the remainder choose the other pure strategy, and that almost every type is strictly optimizing.[14]

We proceed by construction. As a first step, restrict the strategies in $\mathcal{IG}$ so that player 1 must choose either T or B, and player 2 must choose either L or R. Under this restriction, there is an equilibrium of $\mathcal{IG}$ where roughly 2/3 of player 1's types choose T and roughly 2/3 of player 2's types choose R. This equilibrium is determined by a critical type for each player i, namely $\hat{t}_i = \alpha\bar{\varepsilon}$ for α near 1/3, where type t_1 of player 1 chooses

$$\text{B if } t_1 < \hat{t}_1; \text{ and T if } t_1 \geq \hat{t}_1 \tag{2.1}$$

and type t_2 of player 2 chooses

$$\text{L if } t_2 < \hat{t}_2; \text{ and R if } t_2 \geq \hat{t}_2. \tag{2.2}$$

Note that larger types, who assign a higher probability to being found out, choose their favorite pure BoS strategy.

[14] Because Battle of the Sexes has multiple equilibria, so does its associated incomplete information game $\mathcal{IG}$.

For this to be an equilibrium, α must be such that the critical type $\hat{t}_i$ is indifferent between his two pure BoS strategies. From (2.1) and (2.2), the payoff to player 1's critical type $\hat{t}_1 = \alpha\bar{\varepsilon}$ from choosing T is $\alpha 2 + (1-\alpha)0$, if he is not found out, since a fraction α of player 2's types choose L, and 2, if he is found out. Altogether, player $\hat{t}_1$'s payoff from T is

$$\begin{aligned} \pi_{\hat{t}_1}(\mathrm{T}) &= (1-\hat{t}_1)(\alpha 2 + (1-\alpha)0) + 2\hat{t}_1 \\ &= 2(-\bar{\varepsilon}\alpha^2 + (1+\bar{\varepsilon})\alpha). \end{aligned}$$

Similarly, $\hat{t}_1$'s payoff from B is

$$\begin{aligned} \pi_{\hat{t}_1}(\mathrm{B}) &= (1-\hat{t}_1)(\alpha 0 + (1-\alpha)1) + 1\hat{t}_1 \\ &= \bar{\varepsilon}\alpha^2 - \alpha + 1. \end{aligned}$$

These two payoffs are equal when

$$\alpha = \frac{1}{3} - \frac{\sqrt{4\bar{\varepsilon}^2 + 9} - 3}{6\bar{\varepsilon}},$$

which for $\bar{\varepsilon}$ small is close to 1/3, since the second term converges to zero with $\bar{\varepsilon}$. By symmetry, this value of α also makes player 2's critical type $\hat{t}_2 = \alpha\bar{\varepsilon}$ indifferent between L and R.

Now, all types of player 1 below the critical type strictly prefer B to T, whereas all types above strictly prefer T to B. Indeed, for a typical type t_1 of player 1, the difference in payoff from choosing T versus B is

$$\pi_{t_1}(\mathrm{T}) - \pi_{t_1}(\mathrm{B}) = (1-t_1)(3\alpha - 1) + t_1,$$

which, for α close enough to 1/3 (for $\bar{\varepsilon}$ close enough to zero) is strictly increasing in t_1, and vanishes at $\hat{t}_1$. A similar preference holds between L and R for player 2.

Thus the strategies (2.1) and (2.2) form an equilibrium of $\mathcal{IG}$ when the players are restricted to choosing only the pure BoS strategies. We now show that these strategies remain in equilibrium even if this restriction is removed.

In Figure 2.4, the solid lines show player 1's payoff as a function of the probability p he places on T, given that player 2 finds out player 1's mixed strategy. When $p \in [0, 2/3)$, player 2's best reply is R and 1's payoff is decreasing in p. When $p \in (2/3, 1]$, player 2's best reply is L and 1's payoff is increasing in p. When $p = 2/3$, player 2 is indifferent between L and R, but player 1 strictly prefers that player 2 choose L, which accounts for the discontinuity.

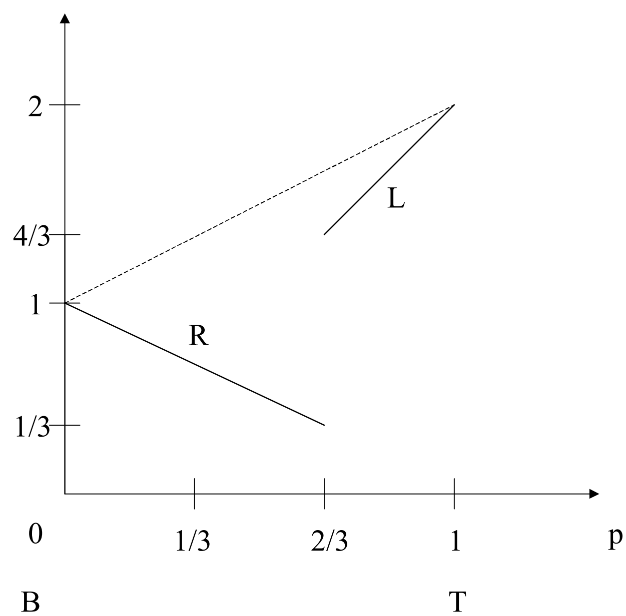

Figure 2.4: Player 1 Found Out in Battle of the Sexes

A mixed strategy in $\mathcal{IG}$ specifies, for each of a player's types, a probability distribution, or "lottery," over the player's mixed strategies in G. If player i of type t_i uses such a lottery, then he assigns probability t_i to the event that his opponent finds out the mixed strategy in G that is the outcome of this lottery.

Consider now the dotted line in Figure 2.4 connecting player 1's payoffs when he chooses the two pure strategies B ($p = 0$) and T ($p = 1$). We claim that, conditional on being found out, any payoff along the dotted line can be achieved by employing an appropriate lottery over the degenerate mixed strategies B and T. In particular, if player 1 of type t_1 uses the lottery that chooses T with probability π and B with probability $1 - \pi$, then player 1's payoff is $2\pi + 1(1 - \pi)$, conditional on being found out. That is, because player 2 finds out the outcome of the lottery, 2's best reply always yields coordination, giving player 1 an average of his payoffs along the diagonal.

The dotted line lies above player 1's payoff in Figure 2.4 and so player 1 prefers such a lottery $\pi \in (0, 1)$ to the mixed strategy giving T the same probability $p = \pi$, conditional on being found out. In contrast to the lottery, the mixed strategy, when combined with the opponent's best reply, leads to miscoordination with positive probability.

On the other hand, conditional on not being found out, the lottery π yields the *same* payoff as does the mixed strategy $p = \pi$. Altogether then, every positive type of player 1 must strictly prefer the lottery $\pi = p$ to the mixed strategy p, for any $p \in (0, 1)$.

Thus, regardless of player 2's strategy, every positive type of player 1 strictly prefers at least one of the two pure strategies T or B to any mixed strategy $p \in (0, 1)$. Consequently, our restriction to these two pure strategies for player 1 is not binding. Similarly, our restriction to player 2's pure strategies L and R is not binding. Hence, the strategies defined in (2.1) and (2.2) are an equilibrium of $\mathcal{IG}$ whose induced distribution on the pure choices converges to the mixed equilibrium of BoS as $\bar{\varepsilon}$ tends to zero.

Further, in this equilibrium, almost every type of each player has a unique best reply. This follows since we showed above that every positive type would be strictly worse off using any non-degenerate mixed strategy, and also that, in the restricted game, every type except the critical one would be strictly worse off using the other pure strategy.

Based upon this equilibrium of $\mathcal{IG}$, our interpretation of the mixed equilibrium of Battle of the Sexes is similar to the Bayesian view: Each player chooses some particular pure strategy, yet the opponent is unsure of which one. The probabilities associated with a player's equilibrium mixture represent the opponent's beliefs about which pure strategy the player will choose.

Thus our incomplete information perturbation is, like Harsanyi (1973), able to rationalize the mixed equilibria of Matching Pennies and Battle of the Sexes as strict equilibria. But the interpretations of the two models are quite distinct. The player types in our perturbation optimally choose whether to reveal or to conceal their choices, choosing to conceal them in Matching Pennies (producing a classical interpretation) and to reveal them in Battle of the Sexes (producing a Bayesian interpretation); whereas in Harsanyi, almost all player types always use only pure strategies.

Our final example leads to a new interpretation of mixed strategy equilibria.

2.3. Modified Rock-Scissors-Paper

Consider the nonzero-sum modification of the zero-sum game Rock-Scissors-Paper shown in Figure 2.5, where $a < b < c < 1$ and a is close to 1. Modified Rock-Scissors-Paper (MRSP) differs from the usual version in two respects. First the

	L	C	R
T	-a,-a	1,-c	-b, 1
M	-c, 1	-b,-b	1,-a
B	1,-b	-a, 1	-c,-c

Figure 2.5: Modified Rock-Scissors-Paper

game is no longer zero-sum because each player receives a payoff near -1 along the diagonal. Second, the off-diagonal payoffs have been perturbed slightly.

The first change adds an element of common interest in that both players now wish to avoid the diagonal. The change in the off-diagonal payoffs is to avoid a particular non genericity, clarified below.[15]

If $a = b = c = 1$, then MRSP has a unique equilibrium in which both players choose each of their pure strategies with probability 1/3. However, because $a < b < c < 1$ and a is near 1, there is a unique equilibrium in which each pure strategy is chosen with probability near 1/3 and in which each player's equilibrium payoff is near $-1/3$.

Figure 2.6 shows player 1's payoff in the incomplete information game $\mathcal{IG}$ as a function of his mixed strategy, conditional on being found out. Triangle TMB in the figure is player 1's simplex of mixed strategy choices. Its vertices are labelled with the pure strategies, T, M and B, they represent. The hyperplanes above the triangle depict player 1's payoff, conditional on player 2 finding out and choosing a best reply. Each hyperplane is labelled with the 2's best reply. The three hyperplanes *almost* meet at player 1's equilibrium strategy in the center of the figure, yielding player 1 a payoff there close to $-1/3$, regardless of 2's best reply.

If player 1 were sure that his strategy would be found out, he would *not* choose a pure strategy, which would result in a payoff of -1; neither would he choose the equilibrium mixture, which yields a payoff near $-1/3$. Instead, player 1 would choose the mixed strategy placing probability 1/2 on T and 1/2 on B. It is then

[15] The remaining coincidences in payoffs are unimportant.

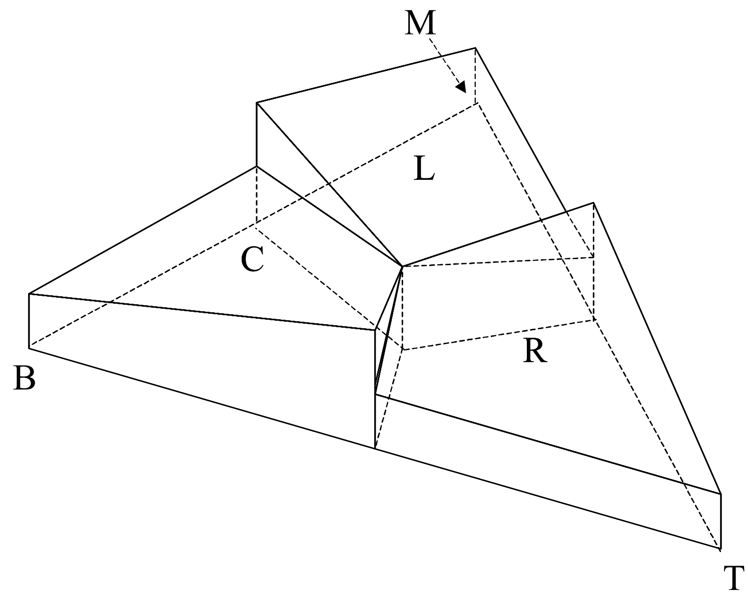

Figure 2.6: Player 1 Found Out in Modified Rock-Scissors-Paper

a best reply for player 2 to choose C resulting in a positive payoff of $(1-a)/2$ for player 1.[16] This achieves coordination in that the diagonal is avoided, but it still conceals 1's final pure choice.

Indeed, the three mixed strategies 1/2-1/2 on T-B; 1/2-1/2 on T-M; 1/2-1/2 on M-B all yield player 1 a positive payoff conditional on being found out. Figure 2.6 shows that these strategies yield attractive payoffs relative to all other mixed strategies, when player 2 chooses a best reply.

If $\bar{\varepsilon}$ is small enough and $\underline{\varepsilon}$ is close enough to $\bar{\varepsilon}$, then these particular three mixed strategies yield an equilibrium of $\mathcal{IG}$, as follows. For each player i there are two critical types, $\hat{t}_{i1} < \hat{t}_{i2}$. Player 1 chooses

$$\begin{aligned} \text{1/2-1/2 on M-B if } t_1 &\in [\underline{\varepsilon}, \hat{t}_{11}) \\ \text{1/2-1/2 on T-M if } t_1 &\in [\hat{t}_{11}, \hat{t}_{12}) \\ \text{1/2-1/2 on T-B if } t_1 &\in [\hat{t}_{12}, \bar{\varepsilon}] \end{aligned}$$

[16]Player 2 is indifferent between C and R, but breaks this tie in player 1's favor by choosing C. Note that player 1, by playing B with a slightly higher probability than T, can give player 2 a strict incentive to choose C.

and player 2 chooses

$$\begin{aligned}
1/2\text{-}1/2 \text{ on C-R if } t_2 &\in [\underline{\varepsilon}, \hat{t}_{21}) \\
1/2\text{-}1/2 \text{ on L-C if } t_2 &\in [\hat{t}_{21}, \hat{t}_{22}) \\
1/2\text{-}1/2 \text{ on L-R if } t_2 &\in [\hat{t}_{22}, \bar{\varepsilon}].
\end{aligned}$$

Moreover, each of these intervals of types occurs with probability approximately 1/3. Each player's strategy therefore induces a probability near 1/3 for each of the original pure strategies, and so approximates the mixed equilibrium of MRSP.

Thus, we are led to the following interpretation of the completely mixed equilibrium of MRSP: Each player i deliberately randomizes by choosing one of the mixed strategies that place probability one-half on each of two pure strategies. The opponent, player j, unaware of which one of the three possible fifty-fifty randomizations player i will employ, assigns probability roughly one-third to each possibility. Player i's equilibrium mixture is obtained by combining the three randomized strategies i might employ according to the weights j's beliefs assign to those strategies.

It is instructive to note the strategic benefits of the above strategies. By choosing a fifty-fifty mixture, enough information is revealed so that, if this mixture is found out, the opponent can successfully avoid the diagonal but cannot take undue advantage. Finally, because a, b, and c are distinct, it can be shown these equilibrium mixed strategies are *strictly optimal.*[17]

We now proceed with the formal analysis of the general case and also explore conditions under which concealment is helpful—as in Matching Pennies and MRSP—and conditions under which it is not—as in Battle of the Sexes.

3. The Incomplete Information Perturbation

Let $G = (u_i, X_i)_{i=1,2}$ be a finite two-person normal form game in which player i's finite pure strategy set is X_i, his mixed strategy set is M_i, and his vNM payoff function is $u_i : X_1 \times X_2 \to \mathbb{R}$. We wish to capture the idea that each player is concerned that the other player might find out his mixed strategy, where the extent to which each player is concerned is private information. Ultimately, we shall be interested in the players' limiting behavior as these concerns vanish.

[17] This is why we perturbed the off-diagonal payoffs.

Given the game G, consider the following associated game of incomplete information.

$$\mathcal{IG} = (U_1, U_2, M_1, M_2, F_1, F_2):$$

- Each F_i is a cdf with $F_i(0) = 0$ and support $T_i \subseteq [0, 1]$.[18]
- Player i's type set is T_i.
- Types are drawn independently according to F_1 and F_2.
- Player i's pure action set is M_i, his set of mixed strategies in G.
- When i's type is t_i and the vector of actions is (m_1, m_2), player i's payoff is

  $$U_i(m_1, m_2, t_i) = (1 - t_i)u_i(m_1, m_2) + t_i v_i(m_i),$$

  where $v_i(m_i)$ is i's payoff in G resulting from m_i together with a best reply against it. If there are multiple best replies for j against m_i, one that is best for i is employed.[19,20]

Thus, $U_i(m_1, m_2, t_i)$ is the payoff i would receive in G when he plays m_i and his opponent plays m_j with probability $1 - t_i$ and plays a best reply to m_i with probability t_i. Player i's type t_i can therefore be interpreted as the probability he assigns to the event that his choice of mixed strategy in G will be found out by the opponent.

The above definition actually yields a collection of incomplete information games indexed by the distribution functions F_1 and F_2. We shall often be concerned with *atomless* cdf's. Such cdf's, F_i, in addition to satisfying $F_i(0) = 0$, are continuous on $[0, 1]$. Note that the incomplete information game approaches the original game G as the measure on each player's types tends to a mass point at zero.

[18] We define the support of a cdf to be the support of the measure it defines, namely the smallest closed set having probability one.

[19] That is, $v_i(m_i) = \max_{x_j} u_i(m_i, x_j)$, s.t. $x_j \in \arg\max_{x'_j \in X_j} u_j(m_i, x'_j)$. So defined, $v_i(\cdot)$ is upper semicontinuous.

[20] This tie-breaking rule is innocuous because generically, some m'_i near m_i leaves j with a unique best reply and gives i a payoff near $v_i(m_i)$.

3.1. Strategies, Lotteries and Induced Distributions

A strategy for player i in $\mathcal{IG}$ is a measurable map from T_i into $\Delta(M_i)$, where $\Delta(M_i)$ denotes the set of Borel probability measures on M_i. We shall refer to elements of M_i as mixed strategies in G, and to elements of $\Delta(M_i)$ as *lotteries on* M_i. So, in the incomplete information game $\mathcal{IG}$, a strategy specifies for each type of each player a lottery over that player's mixed strategies in G. Each player believes that, with the probability given by his type, his opponent finds out the mixed strategy in G that is the outcome of his type's lottery. Pure strategies in the incomplete information game are then degenerate lotteries and so specify a mixed strategy in G for each of a player's types.

An equilibrium of $\mathcal{IG}$ is a pair of strategies that constitute a Nash equilibrium from the ex-ante perspective. Equivalently, an equilibrium strategy pair must be such that given the other player's strategy, the element of $\Delta(M_i)$ chosen by t_i must be optimal for i conditional on t_i, for F_i-almost every t_i.

Let $\sigma_i(\cdot|\cdot)$ be a strategy for player i in $\mathcal{IG}$. Hence, $\sigma_i(\cdot|t_i)$ is for each t_i in T_i a lottery on M_i. Because each m_i in M_i induces a distribution over i's set of pure strategies X_i in G, $\sigma_i(\cdot|\cdot)$ gives x_i in X_i the probability

$$\int_{T_i}\int_{M_i} m_i(x_i)d\sigma_i(m_i|t_i)dF_i(t_i).$$

Let us denote this induced probability by $\bar{\sigma}_i(x_i)$, and the induced mixed strategy in M_i by $\bar{\sigma}_i$.

Because player i's payoff in $\mathcal{IG}$ does not directly depend upon j's type, and because j's strategy σ_j matters to i only when i's strategy is not found out, i's payoff depends only on the induced distribution $\bar{\sigma}_j$ over X_j and not otherwise on σ_j.

This can be seen by considering player i's payoff when his type is t_i and he chooses m_i while his opponent employs the strategy σ_j. Player i's payoff is then

$$(1-t_i)\int_{T_j}\int_{M_j} u_i(m_i,m_j)d\sigma_j(m_j|t_j)dF_j(t_j)+t_iv_i(m_i),$$

which, owing to the linearity of u_i in m_j is equal to

$$(1-t_i)u_i(m_i,\bar{\sigma}_j)+t_iv_i(m_i).$$

4. Zero-Sum Games

In our informal analysis of Matching Pennies in Section 2, we claimed that the equilibrium of $\mathcal{IG}$ in which every type of each player chooses the fifty-fifty mixture is the essentially unique equilibrium. This is a consequence of a more general result for zero-sum games that is given below.

Recall that a maxmin strategy in a zero-sum game is one that yields a player his value if the opponent employs a best reply. We then have the following.

Theorem 4.1. *Suppose that G is a zero-sum game. Then a joint strategy in $\mathcal{IG}$ is an equilibrium if and only if almost every type of each player employs, with probability one, a maxmin strategy for G. Furthermore, in every equilibrium of $\mathcal{IG}$, every type of each player is indifferent among all of his maxmin strategies, and every positive type strictly prefers each of his maxmin strategies to each non maxmin strategy.*

Note that when a player has more than one maxmin strategy, no equilibrium of $\mathcal{IG}$ is strict since all maxmin strategies are then best replies. But the indeterminacy caused by this indifference is inconsequential because any mixture of maxmin strategies is itself a maxmin strategy. Consequently, neither player is forced to employ any particular randomization over his maxmin strategies in order to maintain equilibrium.

Theorem 4.1 leads to a purely classical interpretation of equilibria of two-person zero-sum games, because each player deliberately employs a maxmin strategy (which often involves randomization) and each is certain that the other will do so. We now explore the interpretation our model yields for equilibria of general two-person games.

5. General Two-Person Games

Our objective, as above, is to interpret any equilibrium of G through a nearby equilibrium of $\mathcal{IG}$. To do so requires the game G to be sufficiently robust. The following assumptions make this precise.

5.1. Genericity

Recall from Section 3 that $G = (u_i, X_i)_{i=1,2}$ is a finite two-person game with mixed strategy sets M_i, and that $v_i(m_i)$ is i's payoff in G when he chooses m_i and his opponent plays a best reply to m_i (breaking ties in i's favor if necessary).

For each $x_j \in X_j$, let $C_i(x_j)$ denote those elements of M_i against which x_j is a best reply for j. Consequently, each $C_i(x_j)$ is a convex polyhedron and so possesses finitely many extreme points. Let $\{m_{i1}, ..., m_{iK_i}\}$ denote the union over x_j of the extreme points contained in all the $C_i(x_j)$.

We shall require the following genericity assumptions:

A.1. Every equilibrium of G is regular.

A.2. For each player i, $v_i(m_{i1}), ..., v_i(m_{iK_i})$ are distinct.

Both A.1 and A.2 are satisfied for all but perhaps a closed subset of games, G, having Lebesgue measure zero (in payoff space for any fixed finite number of pure strategies). The proof of this statement is standard in the case of A.1 (see, e.g., van Damme (1991)) and can be found in Appendix B for A.2.

An equilibrium of $\mathcal{IG} = (U_1, U_2, M_1, M_2, F_1, F_2)$ is *essentially strict* if F_i-almost every type of each player i has a unique best choice in M_i. The role of genericity assumption A.2 is to ensure essential strictness, as the following result shows.

Proposition 5.1. *If G satisfies genericity assumption A.2 and each F_i is atomless, then every equilibrium of $\mathcal{IG}$ is essentially strict and almost every type of each player i employs some mixed strategy in $\{m_{i1}, ..., m_{iK_i}\}$.*

The proof can be found in Appendix B. Consequently, for generic games, the problem of indifference does not arise in our incomplete information game. This is important because our player types in general employ non-degenerate mixed strategies. Essential strictness ensures that when non-degenerate mixed strategies are employed, this is not to make the other player indifferent. Rather, they are employed because it is strictly optimal to do so (because concealment happens to be beneficial). We now provide the main result of this section which establishes that our incomplete information game can approximate all equilibria of a generic game G.

5.2. The Main Approximation Theorem

Theorem 5.2. *If G satisfies genericity assumptions A.1 and A.2, then for every $\varepsilon > 0$ there is a $\delta > 0$ such that for all atomless F_1, F_2 satisfying $F_i(\delta) \geq 1 - \delta$ and every equilibrium m^* of G, $\mathcal{IG}$ has an essentially strict equilibrium whose induced distribution on the joint pure strategies in G is within ε of m^*.*

Remark. According to this theorem, for any sufficiently nearby game $\mathcal{IG}$, every equilibrium of G can be approximated by some equilibrium distribution of $\mathcal{IG}$. A standard upper hemicontinuity argument establishes the converse, namely that all equilibrium distributions of nearby games $\mathcal{IG}$ must be close to some equilibrium of G.

The proof is given in Appendix B. The idea is to exploit the fact that player i's behavior in $\mathcal{IG}$ depends only upon the distribution m_j in M_j induced by j's strategy in $\mathcal{IG}$. Moreover, because by A.2 the $v_i(m_{ik})$ are distinct, all but finitely many types of player i have a unique best reply against any such distribution m_j and this best reply is one of the m_{ik}. Letting $g_i(m_j)$ denote the F_i-average over i's best replies as his type varies, it is not difficult to show that g_i is continuous. Moreover, if the mass of F_i is sufficiently concentrated near 0, then $g_i(m_j)$ is very close to a best reply in G against m_j. Consequently, $g = g_1 \times g_2$ is close to the product of the players' best reply correspondences for G. Because any regular equilibrium of G is an "essential" fixed point of G's best-reply correspondence and g is continuous, powerful results from algebraic topology allow us to conclude that g must have a fixed point near any such equilibrium of G. But, by construction, fixed points of g are the distributions on M of equilibria of $\mathcal{IG}$. The desired conclusion follows.

The theorem establishes that our incomplete information perturbation, $\mathcal{IG}$, can rationalize any equilibrium of a generic game G through a nearby equilibrium of $\mathcal{IG}$ in which the players have *strict* incentives to play their part. While this result is reminiscent of Harsanyi (1973), we have already seen that such an equilibrium of $\mathcal{IG}$ sometimes involves a positive measure of a player's types using *non-degenerate* mixed strategies.

5.3. The Interpretation

Let m^* be an equilibrium of a game G that satisfies A.1 and A.2. Suppose that $\mathcal{IG}^n$ converges to G,[21] that σ^n is an equilibrium of $\mathcal{IG}^n$ for every n, and that the induced distributions, $\bar{\sigma}^n$, converge to m^*. By 5.1, almost every type of player i strictly optimizes in σ^n by employing one of the mixed strategies $\{m_{i1}, ..., m_{iK_i}\}$, so that σ_i^n entails some fraction of i's types, μ_{ik}^n, say, employing m_{ik} for each k. Hence, the other players are certain that player i will employ one of the m_{ik}, but they uncertain about *which* of the m_{ik} player i will employ. Their common conjecture, or belief, is that player i will employ m_{ik} with probability μ_{ik}^n. Finally,

[21] That is, the cdf's F_i^n of $\mathcal{IG}^n$ converge to mass points at zero as $n \to \infty$.

because the distribution of σ^n converges to m^*, we must have that for every player i,

$$m_i^* = \mu_{i1}^* m_{i1} + ... + \mu_{iK_i}^* m_{iK_i},$$

where μ_{ik}^* is the limiting fraction of i's types employing m_{ik}.[22] This decomposition of m_i^* therefore leads us to the following interpretation.

Each player i's equilibrium mixture, m_i^*, can be expressed as a convex combination of the mixed strategies $\{m_{i1}, ..., m_{iK_i}\}$. Each mixed strategy given positive weight in the convex combination represents a strategy that i might deliberately employ, while the weight on that mixed strategy represents the opponent's belief that i will employ it.

We have already seen that strictly mixed equilibria in zero-sum games have degenerate decompositions in which all of the weight is placed on the equilibrium mixed strategy. Consequently, such equilibria can always be interpreted from the purely classical point of view where the players deliberately randomize because concealment is beneficial.

Under what conditions is concealment beneficial in the nonzero-sum case? Equivalently, when does the above decomposition place positive weight on at least one non-degenerate mixed strategy? In such cases our interpretation of a mixed equilibrium will involve the classical view. Alternatively, under what conditions will the players instead wish to reveal their pure choices? Equivalently, when does the above decomposition give positive weight only to pure strategies. In such cases, our interpretation is similar to the Bayesian view. (See, for example, The Battle of the Sexes, in Section 2). These questions are taken up next.

6. When To Conceal, When To Reveal

In $\mathcal{IG}$, when a player of a given type strictly prefers to employ a non-degenerate mixed strategy from G, it is because that type strictly prefers concealing the pure choices in the support of that mixed strategy. When this occurs and the equilibrium of $\mathcal{IG}$ is near an equilibrium of G, our interpretation of G's equilibrium will involve (perhaps only partially) the classical view that randomization is deliberate. This motivates the following definition.

[22] Assume without loss that $\mu_{ik}^n \to \mu_{ik}^*$.

Definition 6.1. *An equilibrium m of G is concealing for player i if for some $\eta > 0$ and all sufficiently small $\varepsilon > 0$ there are atomless distributions F_1, F_2 satisfying $F_i(\varepsilon) \geq 1 - \varepsilon$, such that every equilibrium of $\mathcal{IG} \equiv (U_1, U_2; M_1, M_2, F_1, F_2)$ whose distribution on $X_1 \times X_2$ is within ε of m has the property that the F_i-measure of player i's types employing non-degenerate mixed strategies from G is at least η.*

Thus, an equilibrium m of G is concealing if there are nearby incomplete information games in which a positive fraction of types must strictly mix in all equilibria near m.[23] Note that the definition does not require *all* nearby incomplete information games $\mathcal{IG}$ to involve non-degenerate mixing in nearby equilibria. This is appropriate when one takes a more general perspective that includes the present model as well as that of Harsanyi (1973). Because Harsanyi's result establishes that concealment need not play a role in any equilibrium of *some* nearby incomplete information games, the present definition addresses the remaining open question, namely, whether there are *other* nearby incomplete information games in which concealment does play a role.

In fact, however, it may be that, for all nearby incomplete information games of the form $\mathcal{IG}$, all of their equilibria near a given equilibrium of G involve mixing. Consequently, we shall say that an equilibrium, m, of G is *strongly concealing* for player i if it satisfies the definition above when the phrase "there are atomless distributions F_1, F_2 satisfying $F_i(\varepsilon) \geq 1 - \varepsilon$, such that" is replaced by "and for all atomless distributions F_1, F_2 satisfying $F_i(\varepsilon) \geq 1 - \varepsilon$."

Theorem 4.1 implies that a mixed equilibrium of a zero-sum game is concealing for i if and only if it is strongly concealing for i, and this holds if and only if i has no pure maxmin strategy. Consequently, the fifty-fifty equilibrium of Matching Pennies is strongly concealing for both players.

In the nonzero-sum game Modified Rock-Scissors-Paper from Section 2, the unique mixed equilibrium is also strongly concealing for both players. On the other hand, the mixed equilibrium that results when the bottom-left entry in the Battle of the Sexes matrix is changed to (x, x), with $1 < x < 2$, is concealing for both players (which follows from Theorem 6.4 below), but not strongly concealing for either player. As we shall see, the mixed equilibrium of MRSP is strongly concealing because, like the Matching Pennies' equilibrium, its support contains a "cyclic best reply sequence."

[23] Theorem 5.2 ensures that no m can be concealing simply because the particular F_i admit no equilibria of $\mathcal{IG}$ near m.

Formally, a *best reply sequence* in G is a finite sequence $x^1, x^2, ..., x^n$ of joint pure strategies such that in each step, say from x^k to x^{k+1}, one player's strategy is unchanged and the other player's strategy in x^{k+1} is a best reply to the opponent's strategy in x^k. A best reply sequence is *cyclic* if at least two of its elements are distinct and the first and last are identical.

The proofs of the following results can be found in Appendix B.

Theorem 6.2. *Suppose that m^* is an equilibrium of G. If the support of m^* contains a cyclic best reply sequence along which best replies are unique, then m^* is strongly concealing for both players.*

For generic games, players have unique best replies against pure strategies and so along best reply sequences. This leads to the following corollary.

Corollary 6.3. *Generically, if a completely mixed equilibrium is not strongly concealing for either player, then beginning from any joint pure strategy, alternately best replying to one another eventually leads the players to a pure strategy equilibrium.*

Both the theorem and its corollary are driven in part by the fact that when best replies are unique, a best reply sequence can cycle only if somewhere along it some player's payoff strictly falls when the other player switches to a best reply. Intuitively, this makes concealing one's choice attractive. The next result also makes use of this fact.

Theorem 6.4. *Suppose that m^* is a completely mixed equilibrium of G and that each pure strategy of one player is a best reply against some pure strategy of the other. If player i's payoff strictly falls somewhere along a best reply sequence along which best replies are unique, then m^* is concealing for i.*

But suppose that player i's payoff falls nowhere along any best reply sequence. This means that beginning from any joint pure strategy, player i is, generically, made *better off* when player j switches to a best reply against i's strategy. Simply put, player i benefits when j finds out i's pure strategy choice. In such cases one would expect that concealment is harmful, i.e. that player i would prefer to reveal his choice. Our final result shows that this is indeed the case. Note that this result applies, in particular, to The Battle of the Sexes as well as to a whole class of coordination games. In all such games then, our interpretation of their equilibria involves only the Bayesian view. No player deliberately randomizes because randomization is actually harmful.

Theorem 6.5. *Suppose G satisfies genericity assumptions A.1 and A.2. If player i's payoff is weakly increasing along every best reply sequence in G, then no equilibrium of G is concealing for i. Indeed, given A.2, when F_i is atomless, almost every type of player i employs a pure strategy in every equilibrium of $\mathcal{IG}$.*

7. The Decomposition Theorems

We now present two decomposition results. The first shows how our interpretation of any equilibrium of a generic game relates to the solutions of a linear programming problem. The second provides a complete characterization of concealment in terms of such solutions.

Theorem 7.1. *Consider player i's* decomposition problem at $m_i \in M_i$,

$$\max_{\mu_i} \sum_k \mu_{ik} v_i(m_{ik}) \text{ s.t. } \sum_k \mu_{ik} m_{ik} = m_i,$$

where the μ_{ik} must be non negative and sum to one.

Suppose that F_1^n and F_2^n are atomless cdf's whose supports are contained in $[a_n, b_n] \subseteq (0,1)$ and $\lim_n a_n/b_n = 1$. If G satisfies genericity assumption A.2 and σ^n is an equilibrium of $\mathcal{IG}^n = (U_1, U_2, M_1, M_2, F_1^n, F_2^n)$ whose induced distribution converges to m^*, then the vector μ_i^*, whose kth coordinate is the limiting F_i^n-fraction of types t_i who choose m_{ik} according to σ_i^n, is a solution to i's decomposition problem at m_i^*.

Together with Theorem 5.2, Theorem 7.1 implies that under A.1 and A.2, and for every equilibrium m^* of G, at least one solution to i's decomposition problem at m_i^* has a direct interpretation.[24] The weight, μ_{ik}^*, on m_{ik} is the limiting fraction of i's types who employ m_{ik} in a sequence of incomplete information perturbations in which the uncertainty about the players' types tends to zero sufficiently faster

[24] When m_i^* gives no more than three pure strategies positive weight, one can show that i's decomposition problem generically possesses a unique solution. However, there are robust examples with a continuum of solutions, when the support of m_i^* contains four or more pure strategies. Nevertheless, such distinct solutions are always related. In generic games, it can be shown that if μ_i^* solves i's decomposition problem then so does $\hat{\mu}_i$ if and only if for every $x_j \in X_j$, $\mu_i^*(\bar{C}_i(x_j)) = \hat{\mu}_i(\bar{C}_i(x_j))$ and $\sum_{m_{ik} \in \bar{C}_i(x_j)} \mu_{ik}^* m_{ik} = \sum_{m_{ik} \in \bar{C}_i(x_j)} \hat{\mu}_{ik} m_{ik}$, where $\bar{C}_i(x_j) = \{m_{ik} \in C_i(x_j) : v_i(m_{ik}) = u_i(m_{ik}, x_j)\}$. Such indeterminacy is to be expected since i's payoff in $\mathcal{IG}$ is linear on the convex hull of each $\bar{C}_i(x_j)$.

than their concern for being found out. (The significance of limits of this kind is heightened by the comment following the Corollary below.)

The decomposition $m_i^* = \sum_k \mu_{ik}^* m_{ik}$ arising as a solution to i's decomposition problem leads to our interpretation of i's equilibrium strategy m_i^* : Player i employs one of the mixed strategies m_{ik} such that $\mu_{ik}^* > 0$, but player j does not know which one. Player j places probability μ_{ik}^* on the event that player i will employ m_{ik}.

We next turn to a characterization of concealment.

Theorem 7.2. *An equilibrium, m^*, of G is concealing for i if and, under A.1 and A.2, only if every solution to i's decomposition problem at m_i^* places positive weight on a non-degenerate mixed strategy.*

Remark. For a slightly smaller, yet still generic, set of games G, "every solution" can be replaced by "at least one solution" in the statement of Theorem 7.2.

Theorems 7.1 and 7.2 and Definition 6.1 lead directly to the following.

Corollary 7.3. *Suppose that G satisfies A.1 and A.2 and that m^* is an equilibrium of G. Then m^* is concealing for player i if and only if for some $\eta > 0$ and all sufficiently small $\varepsilon > 0$ there exist $0 < a < b < \varepsilon$ such that, for all atomless F_1 and F_2 with supports in $[a, b]$, every equilibrium of $\mathcal{IG} = (U_1, U_2, M_1, M_2, F_1, F_2)$ whose distribution on $X_1 \times X_2$ is within ε of m^* has the property that the F_i-measure of player i's types employing non-degenerate mixed strategies from G is at least η.*

According to the Corollary, whether or not concealment plays a role in a mixed equilibrium can be determined by focussing upon particular sequences of incomplete information perturbations $\mathcal{IG}^n = (U_1, U_2, M_1, M_2, F_1^n, F_2^n)$ converging to G. These are sequences where each F_i^n is atomless with support in $[a_n, b_n]$, and where $a_n, b_n \to 0$ while $a_n/b_n \to 1$.

8. The Significance of Unused Strategies

We now demonstrate that concealment is not an intrinsic property. Whether an equilibrium is concealing, and so whether our interpretation of it involves the classical view, can hinge entirely upon the payoffs to unused strategies. Consider, for example, the game of Figure 8.1. The 2x2 matrix in the top-left corner is

	H	T	U	D
H	1,-1	-1, 1	2, 2	-3,-3
T	-1, 1	1,-1	-3,-3	2, 2
U	2, 2	-3,-3	-3,-3	-3,-3
D	-3,-3	2, 2	-3,-3	-3,-3

Figure 8.1:

Matching Pennies and the Matching Pennies equilibrium remains a regular equilibrium of this game. However, although this equilibrium is strongly concealing in Matching Pennies, without the strategies U and D, it is not even concealing here, when they are present.

To see this, note that both players' payoffs in Figure 8.1 are strictly increasing along every best reply sequence. Because A.1 and A.2 hold generically, we can perturb the game slightly so that both hold and then appeal to Theorem 6.5.

9. $\mathcal{IG}$ as the Reduced Form of a Dynamic Game

Finally, we address a key issue with a static model. This issue was described by von Neumann and Morgenstern as follows (VNM, 17.3, pp.146-8):

> On the one hand we have always insisted that our theory is a static one and that we analyze the course of one play and not that of a sequence of successive plays. But on the other hand we have placed considerations concerning the danger of one's strategy being found out into an absolutely central position. How can the strategy of a player—particularly one who plays a random mixture of several different strategies—be found out if not by repeated observation!

Von Neumann and Morgenstern went on to argue that a dynamic model was nevertheless unnecessary. But their argument is not entirely convincing. It is therefore desirable to sketch a simple dynamic model with the key feature that a player's mixed strategy is sometimes deduced by an opponent who observes the historical frequencies of the player's choices, and such that the static game $\mathcal{IG}$ we have presented in Section 3 captures the dynamic model's equilibrium pattern of play as well as any strict incentives to mix.

Suppose then that the two-person nonzero-sum game, G, is repeatedly played by randomly matching, each period, players from two large populations, so that there is no possibility of two particular players meeting more than once. This familiar device limits the extent to which non-myopic behavior need be considered. Each population is of unit mass, and within each population there are two "varieties" of players. Variety I players, the focus of attention, do not observe the history or variety of an opponent and must pay a positive cost, perhaps arbitrarily small, to implement any strategy that is other than zero-recall.[25] Variety II players are not subject to such a cost and observe an opponent's full history (i.e., the opponent's sequence of past *pure* actions) and his variety before play. Both variety I and variety II players have liminf of the mean payoffs.

Independently of their variety, each player in each population i is assigned a type $t_i \in [0, 1]$ according to the density $h_i(t_i)$. Each player is informed of his variety and his type, both of which are fixed once and for all. Let $\int_0^1 t_j h_j(t_j)dt_j$ be the fraction of variety II players in population i, with the remaining players in population i being variety I players.[26]

Each period every player from each population is matched to a player from the other population as follows. In each population $i = 1, 2$, and for each variety, among the $h_i(t_i)$ players of type t_i a fraction t_i are randomly matched with variety II players from population j and the remaining $(1-t_i)h_i(t_i)$ players are randomly matched with variety I players from population j.[27]

[25] The presence of a complexity cost is merely for simplicity. Appendix C derives the main conclusions obtained here but without the complexity cost.

[26] Each player's type will be the probability that he is matched with a variety II player from the other population. Consequently, the fraction of variety II players in population i must depend upon the density of types in population j, as a matter of accounting.

[27] To see that this is possible, note, for example, that the fraction of variety I players chosen from population 1 to be matched with population 2 variety II players is $\int_0^1 t_1 h_1(t_1)dt_1$. Hence, because there are $\int_0^1 (1-t_2)h_2(t_2)dt_2$ variety I players in population 1, the total number of variety I players chosen from population 1 to be matched with population 2 variety II players is $\int_0^1 t_1 h_1(t_1)dt_1 \int_0^1 (1-t_2)h_2(t_2)dt_2$. A symmetric calculation reveals this to be precisely the

Consequently, a population i player's type, t_i, is the probability that he will be matched with a variety II player from population j. Of course, this is relevant only for variety I players, given that variety II players recognize the variety of their opponent.

Variety I players also need to update the density over their opponent's type after the matching process. Consider then the updated density of a variety I player of any type, t_i say, from population i. Because $(1-t_j)h_j(t_j)$ players of type t_j from population j are randomly matched—independently of their variety—with variety I players in population i, this updated density is independent of t_i and is $f_j(t_j) = (1-t_j)h_j(t_j)t_j / \int_0^1 (1-t_j)h_j(t_j)dt_j$.

Let F_i and H_i denote the distribution functions associated with the densities f_i and h_i, respectively. We claim that if the type distributions satisfy $H_1(\delta) = H_2(\delta) = 1$ for $\delta > 0$ sufficiently small, then any equilibrium, $(m_1^*(\cdot), m_2^*(\cdot))$, of the static game $\mathcal{IG}$ derived from F_1 and F_2 as in Section 3 defines zero-recall strategies used by the variety I's in an equilibrium of the dynamic model. Indeed, suppose that each type t_i of variety I players in population i employs the zero-recall strategy $m_i^*(t_i)$ in the dynamic game and that each variety II player employs a strategy that, in the limit, is guaranteed to choose a one-shot best reply to *any* zero-recall strategy employed by his variety I opponent, breaking ties in this opponent's favor.[28] Suppose also that when two variety II players are matched, they play some one-shot equilibrium of G.

Such strategies are clearly best replies for variety II players in the proposed equilibrium. If variety I players are restricted to zero-recall strategies, each variety

---

number of variety II players chosen from population 2 to be matched with population 1 variety I players. Hence, these two equal sized subpopulations can be randomly matched.

[28] For example, consider the following strategy for each variety II player. Take any sequence $\varepsilon_t > 0$, $\varepsilon_t \to 0$ such that $\varepsilon_t \sqrt{t} \to \infty$, and let m_t denote the empirical frequency of his opponent's choices prior to date t. At date t, the variety II player chooses a pure stage game action so as to maximize the sum of his stage game payoff plus ε_t times his opponent's stage game payoff, given the stage game strategy m_t of his opponent.

To see that this works, suppose a variety I player employs the zero recall strategy, m. Then, by the law of large numbers, $m_t \to m$ almost surely, so that with probability one, all his variety II opponents eventually play, each period, a pure stage game best reply against m. Further, the central limit theorem implies $(m_t - m)/\varepsilon_t \to 0$ with probability one. So, if x and x' are stage game best replies against m for his variety II opponents, but x is strictly better for the variety I player, then the variety II's payoff to x minus the payoff to x', all divided by $\varepsilon_t > 0$, converges with probability one to a positive number. Hence, given their strategy defined above, variety II opponents will, with probability one, eventually never play x' when matched with this variety I player, as desired.

I player's payoff in the dynamic game, given the proposed equilibrium strategies of the others, is precisely described by the payoff function defined in Section 3 for the static model. Hence, in the proposed dynamic game equilibrium, almost every type, t_i, of each variety I player strictly prefers his zero-recall strategy $m_i^*(t_i)$ to any other zero-recall strategy. Can a variety I player gain by deviating to a more complex strategy? If each $H_i(\delta) = 1$ and $\delta > 0$ is sufficiently small, then $m_i^*(t_i)$ is arbitrarily close to a best reply against variety I opponents, and the probability (no greater than δ) of being matched with a variety II opponent is arbitrarily small. Consequently, any such possible gain is outweighed by the implementation cost of the more complex strategy.

Remark. Viewing $\mathcal{IG}$ as a reduced form for the above dynamic model, we see that a player's type, t_i, is not literally the probability that the opponent will find out his mixed strategy. Rather, it is the probability that the opponent has observed the *realizations* of his mixed strategy choices in all *previous* periods. Thus, a player's mixed strategy is never "magically" found out. Rather, it is deduced by an opponent through repeated observations of the player's past actions. The reduced form model, $\mathcal{IG}$, is merely a convenient, parsimonious, representation of this.

Appendices

A. Compatibility of Beliefs

The game $\mathcal{IG}$ described in Section 3 is a perfectly well-defined game of incomplete information. In its most direct interpretation, a player's type is the subjective probability he assigns to the event that his opponent will find out his mixed strategy. But there is a potentially puzzling aspect to this interpretation. While each player believes that the opponent might find out his mixed strategy, evidently neither player considers it possible that he finds out the opponent's mixed strategy. Thus, under this interpretation the players' beliefs appear to be incompatible.[29]

While we do not insist upon the compatibility of the players' beliefs, we now show that the game $\mathcal{IG}$ does not rule out such compatibility. Consider the following extensive form game, where the players know that they will be playing G, but do not necessarily know whether their mixed strategy choices are made simultaneously.

- Nature begins by choosing each t_i, $i = 1, 2$, independently according to H_i on $[0, 1]$.
- Each player i is privately informed of t_i and Nature then determines whether the game is simultaneous according to the following mutually exclusive and exhaustive events.
- With probability $(1 - t_1)(1 - t_2)$ neither player receives any additional information before being required to make a mixed strategy choice (and hence their choices are made simultaneously, i.e., in ignorance of the other's choice).
- With probability $t_i(1 - t_j)$ player i receives no additional information but in fact makes his mixed strategy choice in M_i first, before player j, who is then informed of i's mixed strategy choice prior to choosing a mixed strategy in M_j. (Hence, player i cannot distinguish this event from the previous one.)

[29] This is *not* to say that the beliefs are Harsanyi inconsistent. Indeed, as a purely formal game of incomplete information, $\mathcal{IG} = (U_1, U_2, M_1, M_2, F_1, F_2)$, is entirely standard and involves the common prior $F_1 \times F_2$. It is the *interpretation* of this incomplete information game that raises questions.

- With probability $t_1 t_2$ it is common knowledge (e.g., it is publicly announced) that the two players choose their mixed strategies simultaneously.[30]

- After the players choose their mixed strategies, G is played with those strategies and the game ends.

We now describe how this extensive form game yields $\mathcal{IG}$. Note first that in the extensive form, each player is indeed aware of the possibility that he himself will find out the other player's mixed strategy. Of course, in these subgames, the informed player's decision is entirely straightforward. He simply best replies to the revealed mixed strategy of the opponent. Second, note that when it is common knowledge that the players choose their strategies simultaneously, the resulting game is precisely the simultaneous-move game G, and so any equilibrium of G can be specified in this event regardless of the strategies played in the remaining part of the game. Note also that when the F_i are close to mass points at zero, this common knowledge event has ex-ante probability near zero.

Thus the remaining decision faced by a player i in the above extensive form occurs when he receives no additional information prior to making his mixed strategy choice. In this case, player i assigns probability t_i to the event that the opponent finds out his strategy, just as in $\mathcal{IG}$.

Further, when i receives no additional information, he must update his beliefs concerning j's type (because j's type will typically influence j's choice of mixed strategy).[31] According to Bayes' rule, i's updated beliefs about j's type are given by the distribution

$$F_j(t_j) = \frac{\int_0^{t_j} (1-t) dH_j(t)}{\int_0^1 (1-t) dH_j(t)},$$

which provides the F_1 and F_2 specified in the definition of $\mathcal{IG}$. Notably, this posterior for the opponent's type is independent of own type.[32]

[30] It is not *necessary* to include this "dummy" outcome. Indeed, because it is a common knowledge event, the game tree can be split into two separate trees, in one of which the dummy outcome never occurs. Including this outcome merely simplifies the presentation.

[31] Of course, he can also update his beliefs about the opponent's type when he finds out the opponent's mixed strategy. But such updating is redundant because in this event the opponent's type is of no significance to him.

[32] This independence follows from a simple statistical fact. If X_1, X_2, Y_1 and Y_2 are random variables such that each Y_i depends on X_i, but all other pairs are independent, then conditional

For $i = 1, 2$, let $m_i(t_i)$ denote player i's choice of mixed strategy given his type t_i and when he receives no additional information. Conditional on receiving no information other than his type, player i's ex-ante payoff in the extensive form game is then:

$$\int_0^1 \int_0^1 \{(1 - t_i)u_i(m_1(t_1), m_2(t_2)) + t_i v_i(m_i(t_i))\} dF_1(t_1) \times dF_2(t_2),$$

which is equal to player i's ex-ante payoff in $\mathcal{IG}$ when the players employ the pure strategy pair $(m_1(t_1), m_2(t_2))$. A similar equivalence holds for strategies mapping types into probability distributions (lotteries) on M_i.

So, in the context of this extensive form, $\mathcal{IG}$ focuses attention on the each player's decision in the event that he does not know his opponent's mixed strategy choice and there remains the possibility that the opponent will find out his.

B. Proofs

Proof of Theorem 4.1. The "if" part of the first statement is straightforward. Hence, we proceed with the "only if" part.

Even though G is a zero-sum game, $\mathcal{IG}$ will typically not be. However, $\mathcal{IG}$ is best reply equivalent to the zero-sum game of incomplete information, $\mathcal{IG}^0$, that results when each player i's payoff function is replaced by

$$u_i(m_1, m_2) + \frac{t_i}{1 - t_i} v_i(m_i) - \frac{t_j}{1 - t_j} v_j(m_j).$$[33]

The two games of incomplete information therefore have the same sets of equilibria. Throughout the remainder of the proof, the term "maxmin strategy" will refer to a maxmin strategy in the zero sum game G (not the zero-sum game $\mathcal{IG}^0$).

$\mathcal{IG}^0$ clearly has an equilibrium in which every type of each player chooses a maxmin strategy, giving $\mathcal{IG}^0$ a value of v^0, say.[34] Moreover, because beginning from such an equilibrium player 1's payoff rises above v^0 when a positive measure

on Y_j, X_i provides no addtional information about X_j. Interpreting X_i as the type of player i, and Y_i as the indicator for the event that player j receives additional information yields the independence property described in the text.

[33] This particularly simple argument requires each $\int \frac{t}{1-t} dF_i(t)$ to be finite. A similar proof, which involves a separate argument for types near unity, delivers the result even when one or both integrals are infinite.

[34] If the value of the zero-sum game G is v, then $v^0 = v\{1 + \int \frac{t}{1-t} dF_1(t) - \int \frac{t}{1-t} dF_2(t)\}$.

of player 2's types choose a non-maxmin strategy (owing to the term $-v_2(m_2)$ appearing in 1's payoff and because $F_2(0) = 0$), every equilibrium must involve almost every type of player 2 employing, with probability one, a maxmin strategy. A similar argument applies to player 1. This proves of the "only if" part.

So, in $\mathcal{IG}$, F_j-a.e. type of player j employs one of his maxmin strategies. Consequently, player i can obtain at most his value whether or not he is found out and so is indifferent among all of his maxmin strategies. Furthermore, by employing a non maxmin strategy player i's payoff cannot be above his value if he is not found out and his payoff will be strictly below his value if he is found out. Therefore, every positive type strictly prefers every maxmin strategy to every non maxmin strategy. ■

Proof of Proposition 5.1. The proof relies on two facts. First, for every $m_i \in M_i$, there is a lottery μ_i on $\{m_{i1}, ..., m_{iK_i}\}$ that is at least as good for i as m_i regardless of i's type and regardless of j's strategy. To see this, choose x_j so that $m_i \in C_i(x_j)$ and $v_i(m_i) = u_i(m_i, x_j)$. Clearly, such an x_j exists. Because $m_i \in C_i(x_j)$, m_i is a convex combination of the extreme points of $C_i(x_j)$. We may view the weights in this convex combination as defining a lottery, μ_i, on $\{m_{i1}, ..., m_{iK_i}\}$. Hence, we obtain, for every $t_i \in T_i$ and every strategy σ_j for player j in $\mathcal{IG}$,

$$\begin{aligned}(1-t_i)u_i(m_i, \bar{\sigma}_j) + t_i v_i(m_i) &= (1-t_i)u_i(m_i, \bar{\sigma}_j) + t_i u_i(m_i, x_j) \\ &= \sum_k \mu_{ik} \left[(1-t_i)u_i(m_{ik}, \bar{\sigma}_j) + t_i u_i(m_{ik}, x_j)\right] \\ &\leq \sum_k \mu_{ik} \left[(1-t_i)u_i(m_{ik}, \bar{\sigma}_j) + t_i v_i(m_{ik})\right],\end{aligned}$$

as desired, where the inequality follows because μ_{ik} is positive only when x_j is u_j-best for j against m_{ik} and because, by definition, $v_i(m_{ik}) \geq u_i(m_{ik}, x_j)$ for all such x_j. Note that a consequence of the above inequality is that a player's type has a unique best reply against the opponent's strategy if and only if he has a unique best reply among $\{m_{i1}, ..., m_{iK_i}\}$.

Second, for a fixed mixture $m_i \in M_i$ and a fixed distribution, $\bar{\sigma}_j$, over X_j induced by the opponent's strategy in $\mathcal{IG}$, player i's payoff, $(1-t_i)u_i(m_i, \bar{\sigma}_j) + t_i v_i(m_i)$, is linear in his type t_i. Consequently, because $v_i(\cdot)$ takes on distinct values for distinct extreme points m_{ik}, at most one type can be indifferent between any two of the extreme points.

Together, the two facts imply that at most finitely many types can have multiple best replies among all the extreme points and hence also among all the m_i

in M_i. The result then follows because F_i is atomless. ■

Proof of Genericity of A.2. We wish to show that for fixed finite sets of pure strategies X_1 and X_2, and for all but a closed and Lebesgue measure zero set of pairs of the players' payoff matrices, for each $i = 1, 2$ the values $v_i(m_{i1}), ..., v_i(m_{iK_i})$ are distinct.

Let $n_i = |X_i|$ and let $\mathcal{U}_2$ denote the set of $n_1 \times n_2$ payoff matrices for player 2 in which every submatrix with at least two entries: (i) has full rank after adding a single row of 1's, and (ii) if square, is non singular.

The set of $n_1 \times n_2$ payoff matrices $\mathcal{U}_1$ for player 1 is defined analogously except that "row" is replaced by "column" in (i) above. The usage of "row" and "column" in the following paragraph assumes that $i = 1$ and $j = 2$. In the analogous alternative case, interchange "row" and "column" throughout the paragraph.

Viewing $\mathcal{U}_j$ as a subset of $\mathbb{R}^{n_1 n_2}$, $\mathcal{U}_j$ is open and its complement has Lebesgue measure zero. For any payoff matrix $u_j \in \mathcal{U}_j$, we may construct for each x_j in X_j the convex polyhedral set $C_i(x_j) \subseteq M_i$—which we now write $C_i(x_j; u_j)$ to make explicit the dependence upon u_j. Let $E_i(u_j)$ be the finite union over $x_j \in X_j$ of the finite sets of extreme points of $C_i(x_j; u_j)$. An implication of conditions (i) and (ii) in the definition of $\mathcal{U}_j$ is: $(*)$ if a sequence $u_j^n \in \mathcal{U}_j$ converges to $u_j^0 \in \mathcal{U}_j$, and for every n, m_{i1}^n and m_{i2}^n are distinct elements of $E_i(u_j^n)$ converging to m_{i1}^0 and m_{i2}^0 respectively, then m_{i1}^0 and m_{i2}^0 are distinct elements of $E_i(u_j^0)$. To see this we shall first show that for every $u_j \in \mathcal{U}_j$, $m_i \in E_i(u_j)$ if and only if the submatrix of u_j whose rows are determined by i's pure strategies in the support of m_i and whose columns are determined by j's pure u_j-best replies against m_i, is square.[35] So, suppose first that for some $x_j \in X_j$, m_i is extreme in $C_i(x_j; u_j)$. If the submatrix has fewer rows than columns, then because m_i makes j indifferent between the columns, the submatrix will not have full rank after the addition of a row of 1's, in violation of (i). But if there are fewer columns than rows, then in addition to m_i, there are many linear combinations of the rows, with weights summing to unity, that are proportional to a row of 1's. If $z(x_i)$ denotes the weight on each row x_i in one such solution, z, distinct from m_i, then for $|\alpha| > 0$ small enough $(1-\alpha)m_i + \alpha z$ is in $C_i(x_j; u_j)$ contradicting the fact that m_i is extreme.[36] Hence, the submatrix must be square. Conversely, suppose the submatrix is square. Choose any $x_j \in X_j$ that is u_j-best against m_i. Consequently, m_i is in $C_i(x_j; u_j)$

[35] See also Shapley (1974, Assumption 2.2).

[36] Because $(1-\alpha)m_i + \alpha z$ makes j indifferent between the columns, continuity implies that the columns, and so x_j in particular, are best replies to $(1-\alpha)m_i + \alpha z$ for $|\alpha| > 0$ small enough. Hence, $(1-\alpha)m_i + \alpha z \in C_i(x_j; u_j)$.

and we shall show that m_i is actually extreme in $C_i(x_j; u_j)$. This obviously the case if the submatrix is 1x1, so suppose that it is 2x2 or larger and that m_i is a strict convex combination of distinct elements, m_i', in $C_i(x_j; u_j)$. Each x_j' that is u_j-best against m_i must also be best against each of the m_i', otherwise such an x_j' would not be as good as x_j against m_i. But by (ii) the non singularity of the submatrix implies that, among strategies in M_i—like the m_i'—whose supports are contained in m_i's, m_i is the only one against which each such x_j' is best for j. Thus each $m_i' = m_i$ and we conclude that m_i is extreme in $C_i(x_j; u_j)$. Returning to $(*)$, let us show that $m_{ik}^0 \in E_i(u_j^0)$, $k = 1, 2$. For each $k = 1, 2$, assume without loss that the rows and columns of the submatrix determined by m_{ik}^n and u_j^n are fixed. By the above characterization of the extreme points this submatrix is square and it suffices to show that the submatrix determined by m_{ik}^0 and u_j^0 is square. But the set of rows of the latter submatrix is a subset of those along the sequence because $m_{ik}^n \to m_{ik}^0$, while its set of columns is a superset of those along the sequence because limits of j's best replies remain best replies at the limit. Consequently, the limit matrix has at least as many columns as rows. But it cannot have strictly fewer rows, and so must be square, because m_{ik}^0 makes j indifferent between the columns, and the submatrix would then not have full rank after the addition of a row of 1's, contradicting (i). It remains to show that m_{i1}^0 and m_{i2}^0 are distinct. We have just seen that, for $k = 1$ and 2, the rows and columns determined by m_{ik}^n and u_j^n are the same as those determined by m_{ik}^0 and u_j^0. Thus it suffices to show that the rows and columns determined by m_{i1}^n and u_j^n are not identical to those determined by m_{i2}^n and u_j^n. But this follows immediately from the fact that if they were identical, then the common submatrix they determine is 1x1 or non singular, by (ii), either of which would imply that $m_{i1}^n = m_{i2}^n$, a contradiction. This establishes $(*)$.

For each of player j's payoff matrices $u_j \in \mathcal{U}_j$ define a set of player i's matrices $\mathcal{U}_i(u_j) = \{u_i \in \mathbb{R}^{n_1 n_2} : \sum_{x_i \in X_i} m_{i1}(x_i) u_i(x_i, x_{j1}) \neq \sum_{x_i \in X_i} m_{i2}(x_i) u_i(x_i, x_{j2})$, for all $x_{j1}, x_{j2} \in X_j$ and all $m_{i1} \neq m_{i2}$ s.t. m_{ik} is extreme in $C_i(x_{jk}; u_j)$ for $k = 1, 2\}$. Because X_j is finite and each $C_i(x_j; u_j)$ has finitely many extreme points, $\mathcal{U}_i(u_j)$ is an open subset of $\mathbb{R}^{n_1 n_2}$ whose complement has Lebesgue measure zero.

Let $\mathcal{U}(i) = \{(u_1, u_2) \in \mathbb{R}^{2n_1 n_2} : \ u_j \in \mathcal{U}_j$ and $u_i \in \mathcal{U}_i(u_j)\}$. Note that if $(u_1, u_2) \in \mathcal{U}(1) \cap \mathcal{U}(2)$ then for $i = 1$ and 2, $v_i(m_{ik}) \neq v_i(m_{ik'})$ for all distinct $m_{ik}, m_{ik'}$ in $E_i(u_j)$, as desired. It therefore suffices to show that each $\mathcal{U}(i)$ is open with a Lebesgue measure zero complement. To see that $\mathcal{U}(i)$ is open, suppose that $(u_1^n, u_2^n) \to (u_1^0, u_2^0) \in \mathcal{U}(i)$. Because $\mathcal{U}_j$ is open, u_j^n is eventually in $\mathcal{U}_j$, and $(*)$ implies that u_i^n is eventually in $\mathcal{U}_i(u_j^n)$. Hence, (u_1^n, u_2^n) is eventually in $\mathcal{U}(i)$. To see that the complement of $\mathcal{U}(i)$ has Lebesgue measure zero in $\mathbb{R}^{2n_1n_2}$, note that for every $u_j \in \mathcal{U}_j$, the complement of the section $\mathcal{U}_i(u_j)$ has Lebesgue measure zero in $\mathbb{R}^{n_1n_2}$. Applying Fubini's theorem gives the desired result. ■

The proof of Theorem 5.2 relies on an intuitive corollary of powerful results from algebraic topology.

Corollary B.1. *Suppose U is a bounded, open set in $\mathbb{R}^k$ and f, $g : cl(U) \to \mathbb{R}^k$ are continuous.*[37] *Further, suppose that f is continuously differentiable on U, that x_0 is the only fixed point of f in U, and that $|I - Df(x_0)| \neq 0$. If, for every $t \in [0,1]$, the function $(1-t)f + tg$ has no fixed point on the boundary of U, then g has a fixed point in U.*

Proof of Corollary B.1. Since x_0 is the unique fixed point of f in U, and $|I - Df(x_0)| \neq 0$, it follows that 0 is a regular value of $c(x) = x - f(x)$. Hence, by Dold (1972, IV-5.13.4, p. 71), $\deg_0 c = \text{sgn}|I - Df(x_0)| = \pm 1$. If $d(x) = x - g(x)$, then by hypothesis, for every $t \in [0,1]$, $(1-t)c+td$ has no zero on the boundary of U. Consequently, by Dold (1972, IV-5.13.3, p. 71 and IV-5.4, p. 67), $\deg_0 d = \deg_0 c = \pm 1$ and d has a zero in U. Hence, g has a fixed point in U. ■

Loosely, Corollary B.1 states that if x_0 is the only fixed point of f in some neighborhood, and f is not tangent to the forty-five-degree line, then continuous shifts of f will also have a fixed point in the neighborhood, so long as no fixed point escapes through the neighborhood's boundary.

Proof of Theorem 5.2.[38] Because, by A.1, every equilibrium of G is regular, G has finitely many isolated equilibria. Consequently, it suffices to establish the result for a single equilibrium, m^*, of G. Let $n_i = |X_i|$, and for every $m_i \in M_i$, extend $u_i(m_i, \cdot)$ linearly to all of $\mathbb{R}^{n_j}$. Because, by A.2, the $v_i(m_{ik})$ are distinct for each player i, for every $z_j \in \mathbb{R}^{n_j}$ there is a unique solution, $b_i(z_j|t_i) \in M_i$, to $\max_{m_i \in M_i}(1-t_i)u_i(m_i, z_j) + t_i v_i(m_i)$ for all but perhaps finitely many $t_i \in [0,1]$. Moreover, by the argument given in the proof of Proposition 5.1, the unique maximizer must be one of the m_{ik}. Define $g_i(z_j) = \int_0^1 b_i(z_j|t_i)dF_i(t_i)$. Because F_i is atomless and sufficiently small changes in z_j do not affect the unique best reply of an arbitrarily large fraction of i's types, $g_i : \mathbb{R}^{n_j} \to M_i$ is continuous. Also, note that if $\hat{m}$ is a fixed point of $g = g_1 \times g_2 : \mathbb{R}^{n_1+n_2} \to M$, then $\hat{m} \in M$ and for each player i, $\hat{m}_i = \int_0^1 b_i(\hat{m}_j|t_i)dF_i(t_i)$, so that $(b_1(\hat{m}_2|\cdot), b_2(\hat{m}_1|\cdot))$ is an equilibrium of $\mathcal{IG}$ whose induced distribution on M is $\hat{m}$. Thus, given $\varepsilon > 0$, it

[37] $cl(U)$ denotes the closure of U.

[38] We owe a substantial debt to Hari Govindan who greatly simplified our original proof by providing detailed suggestions upon which the following proof is based.

suffices to show that for all δ small enough, g has a fixed point within ε of m^* whenever $F_i(\delta) \geq 1-\delta$ for $i=1,2$. Henceforth we shall write g^δ to make explicit the dependence of g upon δ.

Because m^* is regular and there are just two players, the number of pure strategies in the support of each player's mixed strategy is the same, l say. So, assume, without loss, that the support of m_i^* is $\{x_{i1}, ..., x_{il}\}$. Define the continuously differentiable function $f_i : \mathbb{R}^{n_1+n_2} \to \mathbb{R}^{n_i}$ by

$$f_i(z_1, z_2) = z_i + \begin{pmatrix} 1 - \sum_k z_{ik} \\ z_{i2}(u_i(x_{i2}, z_j) - u_i(x_{i1}, z_j)) \\ \cdot \\ \cdot \\ z_{in_i}(u_i(x_{in_i}, z_j) - u_i(x_{i1}, z_j)) \end{pmatrix}. \tag{B.1}$$

Let U be an open ball in $\mathbb{R}^{n_1+n_2}$ such that (i) every $z \in U$ is within ε of m^*, the only equilibrium of G in $cl(U)$, (ii) $z \in cl(U)$ implies $z_{ik} > 0$ for every $k \leq l$ and $i = 1,2$, (iii) $z \in cl(U)$ implies $u_i(x_{ik}, z_j) - u_i(x_{i1}, z_j) < 0$ for every $k > l$ and $i = 1,2$. Property (i) can be satisfied because, by regularity, m^* is isolated. Property (ii) can be satisfied because $m_{ik} > 0$ for every $k \leq l$, and property (iii) can be satisfied because x_{i1} is in the support of m_i^* and, by regularity, m^* is quasi-strict.

Remark 1. Letting $f = f_1 \times f_2$, we see that if $\hat{z} \in cl(U)$ is a fixed point of f, then $\hat{z}_{ik}(u_i(x_{ik}, \hat{z}_j) - u_i(x_{i1}, \hat{z}_j)) = 0$ for every $k > 1$, so that by property (ii) of U, $u_i(x_{ik}, \hat{z}_j) - u_i(x_{i1}, \hat{z}_j) = 0$ for every $k \leq l$, and by property (iii) of U, $\hat{z}_{ik} = 0$ for every $k > l$. Also, $\hat{z}$ fixed implies $1 - \sum_k \hat{z}_{ik} = 0$ so that, by property (ii) of U and $\hat{z}_{ik} = 0$ all $k > l$, $\hat{z} \in M$. Consequently, $\hat{z}$ is an equilibrium of G, which means, by property (i) of U, that $\hat{z} = m^*$. Hence, m^*, a fixed point of f in U, is the *only* fixed point of f in $cl(U)$.

Remark 2. Because m^* is regular, $|I - Df(m^*)| \neq 0$, by definition. (See van Damme (1991, p.39).)

Let ∂U denote the boundary of U. We claim that there exists $\bar{\delta} > 0$ small enough such that:

$$\forall \delta < \bar{\delta} \text{ and } \forall t \in [0,1],\ (1-t)f + tg^\delta \text{ has no fixed point in } \partial U. \tag{B.2}$$

Suppose not. Then, because ∂U is compact, there exists $z^\delta \to \hat{z}$, $t^\delta \to \hat{t}$, and $g^\delta(z^\delta) \to \hat{m} \in M$ as $\delta \to 0$ such that for every δ, $(1-t^\delta)f(z^\delta)+t^\delta g^\delta(z^\delta) = z^\delta \in \partial U$. Consequently,

$$(1-\hat{t})f(\hat{z}) + \hat{t}\hat{m} = \hat{z} \in \partial U. \tag{B.3}$$

Furthermore, $\hat{t} > 0$ because otherwise $\hat{z}$ would be a fixed point of f, implying, by Remark 1, that $m^* = \hat{z} \in \partial U$, a contradiction.

Because, for every $\delta > 0$, $g_i^\delta(z_j^\delta)$ is the F_i-average over t_i of maximizers of $(1-t_i)u_i(m_i, z_j) + t_i v_i(m_i)$, and the support of $F_i(\delta) \geq 1-\delta$, $\hat{m}_i$ is a maximizer of $(1-t_i)u_i(m_i, z_j) + t_i v_i(m_i)$ when $t_i = 0$. Hence,

$$\hat{m}_i \text{ solves } \max_{m_i \in M_i} u_i(m_i, z_j). \tag{B.4}$$

So, because, by property (iii) of U, $u_i(x_{ik}, \hat{z}_j) - u_i(x_{i1}, \hat{z}_j) < 0$ for every $k > l$, we must have $\hat{m}_{ik} = 0$ for all $k > l$. Consequently, $\hat{t} > 0$, (B.3) and (B.1) together imply $\hat{z}_{ik} = 0$ for all $k > l$.

If $f_{ik}(\hat{z}) < \hat{z}_{ik}$ for some $1 < k \leq l$, then property (ii) of U and (B.1) imply $u_i(x_{ik}, \hat{z}_j) - u_i(x_{i1}, \hat{z}_j) < 0$ and so by (B.4) $\hat{m}_{ik} = 0 < \hat{z}_{ik}$. But this contradicts (B.3). Consequently, for every $1 < k \leq l$, $f_{ik}(\hat{z}) \geq z_{ik}$ and so by (B.3), and because $\hat{t} > 0$, $\hat{m}_{ik} \leq \hat{z}_{ik}$. Now, if $f_{ik}(\hat{z}) > z_{ik}$ for some $1 < k \leq l$, then property (ii) of U and (B.1) imply $u_i(x_{ik}, \hat{z}_j) - u_i(x_{i1}, \hat{z}_j) > 0$ and so $\hat{m}_{i1} = 0$. By (B.3), this implies $(1-\hat{t})(1-\sum_k \hat{z}_{ik}) + \hat{t}(-\hat{z}_{i1}) = 0$, and because $\hat{z}_{i1} > 0$ by property (ii) of U, we must then have $0 < \hat{t} < 1$ and $1 - \sum_k \hat{z}_{ik} > 0$. But this contradicts $1 = \sum_k \hat{m}_{ik} = \sum_{1<k\leq l} \hat{m}_{ik} \leq \sum_{1<k\leq l} \hat{z}_{ik} < \sum_{k\leq l} \hat{z}_{ik} = \sum_k \hat{z}_{ik}$. Hence, $f_{ik}(\hat{z}) = \hat{z}_{ik}$ for every $1 < k \leq l$, so that by (B.3) and the result of the previous paragraph, $\hat{m}_{ik} = \hat{z}_{ik}$ for all $k > 1$.

Finally, (B.3) implies $(1-\hat{t})(1-\sum_k \hat{z}_{ik}) + \hat{t}(\hat{m}_{i1} - \hat{z}_{i1}) = 0$. But because $\hat{m}_{ik} = \hat{z}_{ik}$ for all $k > 1$ and $\sum_k \hat{m}_{ik} = 1$, we have $1 - \sum_k \hat{z}_{ik} = \hat{m}_{i1} - \hat{z}_{i1}$. Hence, $\hat{m}_{i1} = \hat{z}_{i1}$ and we may conclude that $\hat{z} = \hat{m}$ is a best reply in G against itself. That is, $\hat{m} \in \partial U$ is an equilibrium of G, contradicting property (i) of U, and completing the proof of (B.2).

By (B.2) and Remarks 1 and 2, we may appeal to Corollary B.1 and conclude that for all $\delta < \bar{\delta}$, $g^\delta : \mathbb{R}^{n_1+n_2} \to M$ has a fixed point in U. ■

Proof of Theorem 6.2. Consider the point (x_i', x_j) on the cyclic best reply sequence that maximizes i's payoff when j's pure strategy is a best reply against i's. Consider also the next two points along the sequence, (x_i, x_j) and (x_i, x_j'). Because the sequence is a cycle and best replies are unique along it, $x_i \neq x_i'$ and

$x'_j \neq x_j$. Because the cycle is contained in the support of m^*, $m_i^*(x_i) > 0$ and $m_i^*(x'_i) > 0$.

Now, by construction, $u_i(x'_i, x_j) \geq u_i(x_i, x'_j)$. Also, because best replies are unique along the sequence, $u_i(x'_i, x_j) < u_i(x_i, x_j)$ and we may choose $\gamma > 0$ small enough so that x_j is j's unique best reply against the mixed strategy m_i^γ giving x'_i probability $(1-\gamma)$ and x_i probability γ. Consequently,

$$\begin{aligned} v_i(m_i^\gamma) &= (1-\gamma)u_i(x'_i, x_j) + \gamma u_i(x_i, x_j) \\ &> u_i(x'_i, x_j) \\ &\geq u_i(x_i, x'_j). \end{aligned}$$

But $v_i(x'_i) = u_i(x'_i, x_j)$ and $v_i(x_i) = u_i(x_i, x'_j)$ then imply that $v_i(m_i^\gamma) > v_i(x'_i) \geq v_i(x_i)$. Consequently, m_i^γ is strictly better for i than each of the pure strategies x_i and x'_i when i's strategy is found out.

Suppose σ is an equilibrium of $\mathcal{IG}$. Given the equilibrium strategy σ_j of player j and the distribution, $\bar{\sigma}_j \in M_j$ it induces, suppose without loss that $\min(u_i(x'_i, \bar{\sigma}_j), u_i(x_i, \bar{\sigma}_j)) = u_i(x'_i, \bar{\sigma}_j)$. Then $u_i(m_i^\gamma, \bar{\sigma}_j) = (1-\gamma)u_i(x'_i, \bar{\sigma}_j) + \gamma u_i(x_i, \bar{\sigma}_j) \geq u_i(x'_i, \bar{\sigma}_j)$. Consequently, m_i^γ is at least as good as x'_i when i's strategy is not found out. Altogether, this means that m_i^γ is strictly better than x'_i for every positive type of player i against σ_j.

Consequently, if the distribution $\bar{\sigma}_i$ is close enough to m_i^*, then because the fraction of types employing x'_i is zero and $m_i^*(x'_i) > 0$, a positive and bounded away from zero measure of types must employ non-degenerate mixed strategies. ■

Proof of Theorem 6.4. By Theorem 7.2, it suffices to show that every solution to i's decomposition problem at m_i^* places positive weight on a non-degenerate mixed strategy. So, suppose by way of contradiction that μ_i^* solves i's decomposition problem at m_i^* and that μ_i^* places positive weight only on pure strategies. Because $\sum_k \mu_{ik}^* m_{ik} = m_i^*$, μ_i^* must give each pure strategy x_i weight $m_i^*(x_i)$. Consequently, the optimal value of decomposition problem is

$$\sum_k \mu_{ik}^* v_i(m_{ik}) = \sum_{x_i} m_i^*(x_i) v_i(x_i).$$

By hypothesis there are consecutive elements, $(x_i, x_j), (x_i, x'_j)$, along the best reply sequence such that i's payoff strictly falls from the first to the second. In addition, by hypothesis, x_j is best for j against some x'_i in X_i, and all best replies against pure strategies are unique.

Hence,

$$m_i^*(x_i')v_i(x_i') + m_i^*(x_i)v_i(x_i) = m_i^*(x_i')u_i(x_i', x_j) + m_i^*(x_i)u_i(x_i, x_j'),$$

and $m_i^*(x_i), m_i^*(x_i') > 0$, because the equilibrium is completely mixed.

Choose $\gamma > 0$ small enough so that $\beta_1 \equiv \frac{m_i^*(x_i')}{1-\gamma} \in (0,1)$, $\beta_2 \equiv m_i^*(x_i) - \gamma\beta_1 \in (0,1)$, and x_j is j's unique best reply against the mixed strategy m_i^γ in M_i where $m_i^\gamma(x_i') = (1-\gamma)$ and $m_i^\gamma(x_i) = \gamma$. Hence, $v_i(m_i^\gamma) = u_i(m_i^\gamma, x_j) = (1-\gamma)u_i(x_i', x_j) + \gamma u_i(x_i, x_j)$ and

$$\begin{aligned}
m_i^*(x_i')v_i(x_i') + m_i^*(x_i)v_i(x_i) &= m_i^*(x_i')u_i(x_i', x_j) + m_i^*(x_i)u_i(x_i, x_j') \\
&= \beta_1((1-\gamma)u_i(x_i', x_j) + \gamma u_i(x_i, x_j')) \\
&\quad +\beta_2 u_i(x_i, x_j') \\
&< \beta_1((1-\gamma)u_i(x_i', x_j) + \gamma u_i(x_i, x_j)) \\
&\quad +\beta_2 u_i(x_i, x_j') \\
&= \beta_1 v_i(m_i^\gamma) + \beta_2 v_i(x_i).
\end{aligned}$$

Consequently, because m_i^γ can be written as a convex combination of $m_{ik} \in C_i(x_j)$ and $v_i(m_{ik}) \geq u_i(m_{ik}, x_j)$ for each such m_{ik}, the value of the decomposition problem can be increased by eliminating the weight on x_i' and including instead, for each k, β_1 times the weight on m_{ik} in the convex combination leading to m_i^γ, and decreasing the weight on x_i from $m_i^*(x_i)$ to β_2. Moreover, because $\beta_1 + \beta_2 = m_i^*(x_i') + m_i^*(x_i)$, this improved solution is feasible, contradicting the optimality of the original solution. ■

Proof of Theorem 6.5. Because G satisfies A.1 and A.2, no equilibrium, m^*, of G can be concealing simply because the particular F_i admit no equilibria of $\mathcal{IG}$ near m^*. Consequently, it suffices to prove the second statement.

So, suppose in addition to A.1 and A.2 that F_i is atomless and that σ is an equilibrium of $\mathcal{IG}$. As can be seen from the proof of Proposition 5.1, i's best reply to σ_j, $\sigma_i(t_i)$, is unique and in $\{m_{i1}, ..., m_{iK_i}\}$ for F_i-a.e. t_i. We wish to show that F_i-almost every type of player i employs a pure strategy.

Suppose not. Then a positive F_i-measure of i's types employ some non-degenerate mixed strategy, m_{ik}, say. Suppose that a best reply to m_{ik} for player j which breaks ties in i's favor is x_{jk}. Consider the lottery, μ_i, in $\Delta(M_i)$ giving probability $m_{ik}(x_i)$ to each pure strategy x_i. Regardless of i's type, if j does not find out i's strategy choice, this lottery yields the same payoff as the mixed strategy m_{ik}. If j finds out i's strategy choice, the lottery yields i an expected payoff of

$\sum_{x_i \in X_i} m_{ik}(x_i)v_i(x_i)$, because j finds out the outcome of the lottery. This payoff must be at least as large as $v_i(m_{ik}) = \sum_{x_i \in X_i} m_{ik}(x_i)u_i(x_i, x_{jk})$, since if player j has a best reply to x_i that differs from x_{jk} switching to it cannot hurt player i, by hypothesis. Hence, for each of the positive F_i-measure of types, t_i, for whom m_{ik} is uniquely best against σ_j, we have

$$\begin{aligned} & (1-t_i)u_i(m_{ik}, \bar{\sigma}_j) + t_i v_i(m_{ik}) \\ \leq & \sum_{x_i \in X_i} m_{ik}(x_i)[(1-t_i)u_i(x_i, \bar{\sigma}_j) + t_i v_i(x_i)], \end{aligned}$$

which says that, in equilibrium, t_i's payoff from employing m_{ik} is no higher than his payoff from employing the (distinct) lottery μ_i. But this contradicts the uniqueness of t_i's best reply against σ_j. ■

Proof of Theorem 7.1. By A.2, we may assume that $v_i(m_{i1}) < ... < v_i(m_{iK_i})$, and so by Proposition 5.1, σ^n is essentially strict and F_i^n-a.e. t_i employs some member of $\{m_{i1}, ..., m_{iK_i}\}$. So, for F_i^n-a.e. t_i, $\sigma_i^n(t_i)$ solves

$$\max_k (1-t_i)u_i(m_{ik}, \bar{\sigma}_j^n) + t_i v_i(m_{ik}),$$

which, for all $t_i < 1$ is equivalent to

$$\max_k u_i(m_{ik}, \bar{\sigma}_j^n) + \frac{t_i}{1-t_i} v_i(m_{ik}).$$

Now, because $v_i(m_{ik})$ is strictly increasing in k, types choosing m_{ik} will be lower than those choosing $m_{ik'}$ for $k' > k$. Consequently, σ_i^n specifies for each m_{ik} an interval of (possibly degenerate) types, $[t_{ik-1}, t_{ik}]$, employing that strategy. Letting μ_{ik}^n denote the F_i^n-fraction of types t_i who employ m_{ik}, the μ_{ik}^n sum to unity over k, the distribution induced by σ_i^n is $\bar{\sigma}_i^n = \sum_k \mu_{ik}^n m_{ik}$, and μ_{ik}^n, $k = 1, ..., K_i$ must solve

$$\max_{\{\mu_{ik}, t_{ik}\}} \sum_k \mu_{ik} u_i(m_{ik}, \bar{\sigma}_j^n) + \sum_k E_{F_i^n}\left(\frac{t_i}{1-t_i} \middle| t_i \in [t_{ik-1}, t_{ik}] \right) \mu_{ik} v_i(m_{ik}),$$

subject to $\sum_k \mu_{ik} = 1$, $\mu_{ik} \geq 0$, and $F_i^n(t_{ik}) - F_i^n(t_{ik-1}) = \mu_{ik}$ for every k, where $t_{i0} = 0$.

But because the first sum reduces to $u_i(\bar{\sigma}_i^n, \bar{\sigma}_j^n)$ at the optimum, μ_{ik}^n, $k = 1, ..., K_i$ must also solve

$$\max_{\{\mu_{ik}, t_{ik}\}} \sum_k E_{F_i^n}\left(\frac{t_i}{1-t_i} \middle| t_i \in [t_{ik-1}, t_{ik}] \right) \mu_{ik} v_i(m_{ik}),$$

subject to $\sum_k \mu_{ik} m_{ik} = \bar{\sigma}_i^n$, $\sum_k \mu_{ik} = 1$, $\mu_{ik} \geq 0$, and $F_i^n(t_{ik}) - F_i^n(t_{ik-1}) = \mu_{ik}$ for every k, where $t_{i0} = 0$.

Now, because the support of F_i^n is contained in $[a_n, b_n]$ and $\lim a_n/b_n = 1$, $\frac{1}{b_n} E_{F_i^n}\left(\frac{t_i}{1-t_i}\middle| t_i \in [\alpha_n, \beta_n]\right) \to 1$ as $n \to \infty$ whenever $F_i^n(\beta_n) - F_i^n(\alpha_n) > 0$, all n. Consequently, premultiplying the above maximand by $1/b_n$, taking the limit as $n \to \infty$, and applying Berge's Theorem of the Maximum (the constraint correspondence being continuous) implies that μ_i^* solves

$$\max_{\{\mu_{ik}\}} \sum_{k=1}^{K_i} \mu_{ik} v_i(m_{ik}) \text{ s.t. } \sum_{k=1}^{K_i} \mu_{ik} m_{ik} = m_i^*,$$

where the μ_{ik} are non negative and sum to one over k. ■

Proof of Theorem 7.2. Suppose that every solution to i's decomposition problem at m_i^* places positive probability on some non-degenerate mixed strategy. Let F_1^n and F_2^n be atomless cdf's whose supports are contained in $[a_n, b_n]$, where $a_n/b_n \to 1$ as $n \to \infty$. We claim that if σ^n is an equilibrium of $\mathcal{IG}$ when the cdf's are F_i^n, and $\bar{\sigma}^n$ converges to m^*, then the limiting F_i^n-fraction of types who employ some non-degenerate mixed strategy is strictly positive. Hence, m^* is concealing.

To see this, note first that even though we have not assumed A.2, the proof of Proposition 5.1 nonetheless shows that for every m_i, there is a lottery on $\{m_{i1}, ..., m_{iK_i}\}$ that induces the distribution m_i and is at least as good as m_i for every t_i regardless of j's strategy σ_j. Consequently, we may replace $\sigma_i^n(t_i)$ with a payoff- equivalent and distribution-equivalent lottery on $\{m_{i1}, ..., m_{iK_i}\}$.[39] Moreover, the limiting F_i^n-fraction of types who employ some non-degenerate mixed strategy is strictly positive according to the newly defined equilibrium strategy if and only if this is so according to σ^n. Thus we may assume, without loss, that each $\sigma_i^n(t_i)$ is a lottery on $\{m_{i1}, ..., m_{iK_i}\}$. It now suffices to show that the probability vector, μ_i, whose kth coordinate is the limiting F_i^n-fraction of types who employ m_{ik} is a solution to i's decomposition problem. But this now follows exactly as in the proof of Theorem 7.1.

Conversely, suppose that m^* is concealing and that A.1 and A.2 hold. Then there is an $\eta > 0$, a sequence $\varepsilon_n \to 0^+$ and atomless distributions F_1^n, F_2^n satisfying $F_i(\varepsilon_n) \geq 1 - \varepsilon_n$ such that every equilibrium of $\mathcal{IG}^n \equiv (U_1, U_2; M_1, M_2, F_1^n, F_2^n)$ whose distribution on X is within ε_n of m^* has the property that the F_i^n-measure

[39] Measurability can also be preserved by ensuring that the map from M_i to lotteries on $\{m_{i1}, ..., m_{iK_i}\}$ is measurable.

of player i's types employing non-degenerate mixed strategies from G is at least η.

By Theorem 5.2, there is a corresponding sequence, σ^n, of essentially strict equilibria of $\mathcal{IG}^n$ whose distributions on M converge to m^* as n tends to infinity. Let μ^n_{ik} denote the measure of player i's types employing m_{ik} in the equilibrium σ^n. Because the equilibrium is essentially strict, the μ^n_{ik} sum to unity over $k = 1, .., K_i$. Moreover, the sum of the μ^n_{ik} over those k corresponding to non-degenerate mixed strategies m_{ik} is at least $\eta > 0$ for every n. Consequently, letting $\mu^*_{ik} = \lim_n \mu^n_{ik}$—convergence may be assumed without loss—it must be the case that μ^*_i places a total weight of at least $\eta > 0$ on non-degenerate mixed strategies m_{ik}. Moreover, because the distribution of σ^n converges to m^* we must have $\sum_k \mu^*_{ik} m_{ik} = m^*_i$, so that μ^*_i is feasible for i's decomposition problem.

Now, for every $\mu^*_{ik} > 0$ such that m_{ik} is non-degenerate, it must be the case that $\sum_{x_i \in X_i} m_{ik}(x_i) v_i(x_i) < v_i(m_{ik})$. (Equality trivially holds if m_{ik} is degenerate.) This is because the opposite inequality would imply that the positive measure of types employing m_{ik} in each equilibrium σ^n are not strictly optimizing—the lottery giving probability $m_{ik}(x_i)$ to x_i is at least as good as m_{ik}—contradicting the essential strictness of the equilibrium.

Consequently, $\sum_{x_i \in X_i} m^*_i(x_i) v_i(x_i) = \sum_{x_i \in X_i} \left(\sum_k \mu^*_{ik} m_{ik}(x_i)\right) v_i(x_i) < \sum_k \mu^*_{ik} v_i(m_{ik})$. But this means that *every* solution to i's decomposition problem must place strictly positive weight on some non-degenerate mixed strategy, since the only way to obtain m^*_i as a convex combination of only pure strategies is to employ the weights $m^*_i(x_i)$, and this evidently yields a suboptimal value for the decomposition problem. ■

C. Eliminating the Complexity Cost in the Dynamic Game

Consider the dynamic game from Section 9, but without the complexity cost. The equilibrium described there need no longer be an equilibrium. However, this can be remedied by modifying only the strategy of the variety II players. Recall that $(m^*_1(\cdot), m^*_2(\cdot))$ is an equilibrium of the static game $\mathcal{IG}$. We wish to show that in the dynamic game it is an equilibrium for population i variety I players of type t_i to employ the zero-recall strategy $m^*_i(t_i)$ in every period and for variety II players in both populations to employ the strategy, $\hat{s}_{II}$ defined as follows. As a first step, define the following variety II strategy, denoted s^0_{II}. (This modifies that in Footnote 28.) For $\varepsilon_t = 1/\log t$, the variety II player chooses, at date t, a pure stage game action, x, to maximize, against the empirical distribution

m_t of the opponent, his stage game payoff plus ε_t times the opponent's stage game payoff plus ε_t^2 times $\phi(x)$, where ϕ takes on distinct values for distinct stage game strategies. Set $\bar{m}_j = \int m_j^*(t_j)dF_j(t_j)$. As a second step, for $i = 1, 2$ and $t_i \in [0, 1]$, consider the two-person zero-sum game, $Z_i(t_i)$, between players 1 and 2, where the payoff to i is $(1 - t_i)u_i(m_i, \bar{m}_j) + t_i u_i(m_i, m_j)$, when i and j employ $m_i \in M_i$ and $m_j \in M_j$ respectively. Let $\tilde{m}_j(t_i)$ be an optimal strategy for j in this zero-sum game, and, for $i = 1, 2$, let $\{t_{i1}, t_{i2}, ...\}$ be a countable dense subset of $[0, 1]$. We can now define a suitable variety II strategy, $\hat{s}_{II}$, as follows. At any date t between $n!$ and $(n + 1)!$, and for every t' and t'' between $(n - 1)!$ and $n!$, compare the empirical distributions, $m_{t'}$ and $m_{t''}$, of the actions of the variety I player to whom the variety II player is currently matched. If $\|m_{t'} - m_{t''}\| < 1/n^6$ for every such t' and t'' (as would almost surely eventually be the case, by the central limit theorem, if the variety I player employed a zero-recall strategy) then the variety II player plays according to s^0_{II}, and we say that the interval of dates $[(n - 1)!, n!]$ "passes the Cauchy test." Otherwise, the variety II player, if he is from population j, plays $\tilde{m}_j(t_{ik})$, where k is the number of $m \in \{1, 2, ..., n\}$ such that $[(m-1)!, m!]$ does not pass the Cauchy test. Given these strategies, note that, because the Cauchy test is eventually always passed when variety I players use zero-recall strategies, the variety II players are optimizing and $m_i^*(t_i)$ is the best zero-recall strategy for a population i variety I player of type t_i. It remains to show that no variety I player can profitably deviate to a supergame strategy. Suppose then that one particular variety I player, hereafter "the deviator," deviates to some supergame strategy. Since, in every period, the stage game mixture employed by any player matched against the deviator depends at most upon the deviator's own past actions, the deviator's payoff from any deviation can be obtained with a (possibly mixed) supergame strategy whose current period behavior depends only upon his own past actions.[40] Hence, if there is a profitable deviation, there is a profitable pure supergame strategy deviation depending only upon the deviator's own past actions. Any such strategy generates a fixed infinite sequence of his actions, regardless of the others' actions. Hence it

[40] Consider any repeated game. If all opponents of player i employ date t pure strategies, s^t_{-i}, depending only on i's past actions, $a_i^0, a_i^1, ..., a_i^{t-1}$, and i employs the unrestricted pure strategy $s_i = (s_i^t)_{t=0}^{\infty}$, then define the strategy $\hat{s}_i$ inductively by $\hat{s}_i^0 = s_i^0$, $\hat{s}_i^{t+1}(a_i^0, a_i^1, ..., a_i^t) = s_i^{t+1}(a_i^0, a_i^1, ..., a_i^t, s^0_{-i}, s^1_{-i}(a_i^0), ..., s^t_{-i}(a_i^{t-1}))$. The strategy $\hat{s}_i$ then depends only upon i's past actions and the strategy pairs (s_i, s_{-i}) and $(\hat{s}_i, s_{-i})$ yield the same infinite sequence of actions. When s_i and s_{-i} are mixed, a similar construction yields a mixed strategy $\hat{s}_i$ depending only on i's past actions such that (s_i, s_{-i}) and $(\hat{s}_i, s_{-i})$ induce the same probability distributions over infinite sequences of actions.

suffices to show that the deviator cannot improve his payoff by deciding at the start of the game to switch to some fixed infinite sequence of actions. Now, any such sequence either (i) generates a sequence of empirical distributions of actions such that, for all large enough n, the interval $[(n-1)!, n!]$ passes the Cauchy test or (ii) it does not. In case (i), the deviator's sequence of empirical distributions, m_t, satisfies $\|m_t - m_r\| < \frac{1}{n^6} + \frac{1}{(n+1)^6} + ... + \frac{1}{(n+k+1)^6}$ if t is between $(n-1)!$ and $n!$ and r is between $(n+k)!$ and $(n+k+1)!$. Hence, this sequence is indeed Cauchy, converging to some $\hat{m}$ which, taking the limit as $r \to \infty$ in the previous inequality, satisfies $\|m_t - \hat{m}\| \leq \frac{1}{n^6} + \frac{1}{(n+1)^6} + ... < \infty$ for t between $(n-1)!$ and $n!$. Consequently, for such t and n, $\|m_t - \hat{m}\| / \varepsilon_t^2 = (\log t)^2 \|m_t - \hat{m}\| \leq (\log n!)^2 \|m_t - \hat{m}\| \leq (\frac{n(n+1)}{2})^2 (\frac{1}{n^6} + \frac{1}{(n+1)^6} + ...) \to 0$ as $t, n \to \infty$, so that his variety II opponents (whenever matched against him) eventually play the unique constant pure action equal to their best reply against $\hat{m}$ that maximizes the deviator's payoff, breaking any double-ties according to ϕ. Hence, the deviating variety I player's payoff from this particular sequence of actions is the same as his payoff from the zero-recall strategy $\hat{m}$, which can be no higher than his payoff from his best zero-recall strategy, a payoff that can be obtained without deviating. In case (ii), the Cauchy test fails infinitely often and so if the deviating variety I player is from population i, he induces his variety II opponents (whenever so matched) to play $\tilde{m}_j(t_{ik})$ in every period t between dates $n_k!$ and $(n_k+1)!$ for some $n_k \to \infty$, where $k = 1, 2, ...$ is the number of Cauchy test failures prior to date n_k. Whatever his type t_i, there is a subsequence $t_{ik_l} \to t_i$ as $l \to \infty$. Hence, $\tilde{m}_j(t_{ik_l}) \to \tilde{m}_j(t_i)$ and so, when matched against him, his variety II opponents play stage game strategies that converge to $\tilde{m}_j(t_i)$ in every period t between dates $n!$ and $(n+1)!$, along some subsequence of $n = 1, 2, ...$. Given the behavior of his variety I opponents, his expected payoff in every period t between such dates $n!$ and $(n+1)!$ is then, in the limit, no larger than his value in the zero-sum game $Z_i(t_i)$. Hence the liminf of his expected mean payoff can be no larger than his value in $Z_i(t_i)$. By Fatou's lemma, it follows that his payoff in the dynamic game, namely the expected liminf of the mean payoff, is then no greater than his payoff in the equilibrium $(m_1^*(\cdot), m_2^*(\cdot))$ of the static game $\mathcal{IG}$. However, he can obtain this latter payoff in the dynamic game by playing the zero-recall strategy $m_i^*(t_i)$ and so the deviation is not profitable in this case either. Therefore, his type's best zero-recall strategy, $m_i^*(t_i)$, is optimal among all supergame strategies. That this zero-recall strategy might be strictly mixed and strictly better than any other pure zero-recall strategy follows as before.

References

Armbruster, W., and W. Boege (1979): "Bayesian Game Theory," in *Game Theory and Related Topics*, ed. by O. Moeschin and D. Pallaschke. Amsterdam: North-Holland.

Aumann, R.J., Y. Katznelson, R. Radner, R. Rosenthal, and B. Weiss (1983): *Mathematics of Operations Research*, 8, 327-341.

Aumann (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, pp.1-18.

Aumann and Brandenburger (1995): "Epistemic Conditions for Nash Equilibrium," *Econometrica,* 63, pp.1161-1180.

Brandenburger, A. and E. Dekel (1989): "The Role of Common Knowledge Assumptions in Game Theory," in *The Economics of Missing Markets, Information, and Games,* ed. by F. Hahn. Oxford: Oxford University Press, 46-61.

Dold, A. (1972): *Lectures on Algebraic Topology.* Berlin: Springer.

Harsanyi, J.C. (1973): "Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points," *International Journal of Game Theory*, 2, 1-23.

Matsui, A. (1989): "Information Leakage Forces Cooperation," *Games and Economic Behavior,* 1, 94-115.

Radner, R. and R. Rosenthal (1982): "Private Information and Pure-Strategy Equilibria," *Mathematics of Operations Research*, 7, 401-409.

Robson, A.J. (1994): "An 'Informationally Robust Equilibrium' for Two-Person Nonzero-Sum Games," *Games and Economic Behavior,* 7, 233-245.

Rockafellar, R.T. (1970): *Convex Analysis.* Princeton: Princeton University Press.

Rosenthal, R.W. (1991): "A Note on Robustness of Equilibria with Respect to Commitment Opportunities," *Games and Economic Behavior,* 3, 237-243.

Schelling, T.C. (1960): *The Strategy of Conflict.* Cambridge, MA: Harvard University Press.

Shapley, L.S. (1974): "A Note on the Lemke-Howson Algorithm," *Mathematical Programming Study,* 1, 175-189.

Tan, T. and S. Werlang (1988): "The Bayesian Foundations of Solution Concepts of Games," *Journal of Economic Theory*, 45, 370-391.

Van Damme, E. (1991): *Stability and Perfection of Nash Equilibria.* Berlin: Springer-Verlag.

Von Neumann, J. and Morgenstern, O. (1944): *Theory of Games and Economic Behavior.* Princeton: Princeton University Press.