

Payoff Information and Self-Confirming Equilibrium¹

Eddie Dekel

*Department of Economics, Northwestern University, Evanston, Illinois 60201 and
Department of Economics, Tel Aviv University, Tel Aviv, Israel*

Drew Fudenberg

Harvard University, Cambridge, Massachusetts 02138

and

David K. Levine

*Department of Economics, University of California at Los Angeles,
Los Angeles, California 90095-1477*

Received August 5, 1997; revised July 20, 1999

In a self-confirming equilibrium, each player correctly forecasts the actions that opponents will take along the equilibrium path, but may be mistaken about the way that opponents would respond to deviations. This paper develops a refinement of self-confirming equilibrium in which players use information about opponents' payoffs in forming beliefs about the way that opponents play off of the equilibrium path. We show that this concept is robust to payoff uncertainty. We also discuss its relationship to other concepts and show that it is closely related to assuming almost common certainty of payoffs in an epistemic model with independent beliefs. *Journal of Economic Literature* Classification Numbers: C72, D84. © 1999 Academic Press

1. INTRODUCTION

Suppose, as is now common, that we interpret equilibrium in a game as a steady state of some non-equilibrium process of adjustment and “learning.” What steady states might we expect to observe, and conversely what strategy profiles seem unlikely to be steady states? The notion of

¹ We thank Farak Gul, an associate editor, two referees, and seminar participants at Carlos III University, Stanford GSB, UCLA, and University College, London for helpful comments. Financial support from the UCLA Academic Senate, and NSF Grants SBR-93-20695, SBR-9409180, and SBR-94-09302 are gratefully acknowledged. A portion of this research was conducted while Dekel and Levine were visiting the School of Economics Tel-Aviv University.

self-confirming equilibrium is designed to model steady states where players have no *a priori* information about opponents' play or payoffs, and each time the game is played they observe only the actions played by their opponents. Intuitively, self-confirming equilibrium requires only that players correctly forecast the actions opponents will take along the equilibrium path, but does not require that their off-path beliefs are correct. When the only information players have is the observed play in the game, they will never receive evidence that their forecasts of off-path play are incorrect. We expect, then, that any self-confirming equilibrium can be a steady state, including those with outcomes that cannot arise in Nash equilibrium.²

Because self-confirming equilibrium (henceforth "SCE") allows beliefs about off-path play to be completely arbitrary, it does not force the beliefs to incorporate restrictions that players might be able to deduce from information about opponents' payoff functions. That is, it supposes that players do not "think strategically," but simply learn from their experience. Of course, if players have no information about their opponents' preferences, they are unable to deduce that the opponents like certain actions more than others, and there is no reason to restrict beliefs about off-path play. This may be a good approximation of some real-world situations, and is also the obvious way to model play in experiments in which subjects are given no information about opponents' payoffs. In other cases, both in the real world and in the laboratory, it seems plausible that players have and use some information about their opponents' payoffs.³

The goal of this paper is to develop a more restrictive version of SCE that captures how players' deductions, based on commonly known information about all the payoffs in the game, can restrict the set of observed long-run outcomes. We provide some formal results to verify the robustness of our solution concept to certain perturbations, and to help relate our contribution to past work, but these results are not the main point of the paper. Rather, the paper's main contribution is the development of the "rationalizable self-confirming" concept, and the illustration of its implications in a number of examples.

² Fudenberg and Kreps [10] and Fudenberg and Levine [13].

³ In many experiments subjects are told the rules that determine their opponents' money payoffs. The extent to which this approximates common knowledge of payoffs depends on the extent to which opponents are believed to be motivated by non-pecuniary factors such as altruism or spite. In some experiments, there is evidence that a substantial fraction of subjects are motivated by non-pecuniary factors. But there is also experimental evidence that some players successfully apply concepts such as iterated dominance to anticipate opponents play; see, for example, Costa-Gomez *et al.* [8]. However, there is substantially more scope for the experimental study of the impact of information about other players' payoffs. This paper suggests the hypothesis that without payoff information we should expect to see an SCE, but with the additional information, we should see only RSCE.

The key issue is how to incorporate the information about payoffs into SCE. We suppose that players believe that their opponents' actions will maximize their presumed payoff functions *so long as the opponents have not been observed to deviate from anticipated play*. However, players do not use the prior payoff information to restrict their beliefs about the play of opponents who have already been observed to deviate from expected play. Intuitively, this corresponds to players supposing that such deviations are signals that the deviator's payoff function is different than had been expected. More formally, we require that a player's strategy be optimal at all of his information sets that are not precluded by the strategy itself; we call these *reachable* information sets.⁴

There are two related reasons that we impose optimality only at *reachable* off-path information sets, rather than at all information sets. In Section 2 we show by example that the latter requirement is not robust in the sense of Fudenberg *et al.* [11], and in Section 4 we prove that the former is robust. In Section 5 we claim that optimality at reachable information sets follows from a natural epistemic model that assumes caution and almost common knowledge (in the sense of Monderer and Samet [19]) of rationality.⁵ Reny [21, 22], Ben Porath [3], and Gul [16], among others, also argue (in varying degrees of specificity) for optimality at reachable nodes.

To capture the idea that play corresponds to the steady state of a learning process in which the path is observed each time the game is played, we also assume that the path of play is public information. This is in the spirit of, but stronger than, the assumption underlying self-confirming equilibrium, which is that each player knows the path of play. For simplicity, we also impose the assumption that players' beliefs concerning their opponents' play correspond to independent randomizations. Combining these assumptions with that of optimality at reachable nodes leads to "rationalizable self-confirming equilibrium," or "RSCE."

Papers by Rubinstein and Wolinsky [23] and Greenberg [15], like ours, are based on the idea that players form their forecasts of opponents play using both prior information both about the opponents' payoffs and

⁴ There are two closely related notions of optimality at off-path information sets that we consider: best replies to the *limit* of a sequence of trembles, namely sequential rationality, as in Kreps and Wilson [17], and best replies to the sequence itself, as in Selten's notion [24] of trembling-hand perfection. We expect that, as in the relationship between sequential and perfect equilibrium, the difference is only in non-generic games—see Kreps and Wilson [17] and Blume and Zame [6]—but verifying this takes us too far afield. In this introduction we are imprecise and use optimality to refer to both notions.

⁵ The relationship between the results of Sections 4 and 5 is similar to that between the results of Dekel and Fudenberg [9] and those of Börgers [7].

some information about what is observed when the game is actually played. Greenberg's notions of null mutually acceptable courses of action and path mutually acceptable courses of action correspond to the non-robust "sequentially rationalizable sets" and "sequentially rationalizable self-confirming equilibrium" that we define in Section 4. However, where we use these concepts only as tools for understanding the RSCE concept, Greenberg uses them as the center of his analysis.

Rubinstein and Wolinsky consider strategic-form, instead of extensive-form games and therefore do not impose optimality at off-path information sets. They represent the information that players obtain about opponents' play by arbitrary but deterministic "signal functions" and they allow for correlated beliefs. These differences complicate the comparison of their work with ours, and we defer a fuller discussion to the Section 5, but roughly speaking, in the cases that are common to their work and ours, their notion of an RCE ("rationalizable conjectural equilibrium") corresponds to self-confirming equilibrium—the "rationalizable" aspect of their concept has no additional bite.

We should make clear from the outset that, although this paper is motivated by the learning-theoretic approach to equilibrium in games, we do not here provide an explicit learning-theoretic foundation for our concepts. We are confident that such foundations can be constructed by, for example, incorporating restrictions on the priors into the steady-state learning model of Fudenberg and Levine [13], but we have not checked the details.

2. THE SOLUTION CONCEPTS

There are n players $i = 1, \dots, n$ and a game tree with decision nodes $x \in X$ and terminal nodes $z \in Z$. Information sets for player i are $h_i \in H_i$; singleton information sets for nature are H_0 . Available actions at an information set are $A(h_i)$. A behavior strategy for player i , $\pi_i \in \Pi_i$, satisfies $\pi_i(h_i) \in A(A(h_i))$. The payoffs are $u_i(z)$.

A profile π is said to have the same path as $\hat{\pi}$ if $\pi, \hat{\pi}$ agree on the set of information sets reached with positive probability under π . Call an information set h_i *reachable* given strategy π_i if there exists a π'_{-i} such that h_i is reached with positive probability given (π_i, π'_{-i}) ; These represent places in the tree that are consistent with player i playing π_i .

An *assessment* a_i for player i is a probability distribution over nodes at each of his information sets. A *belief* for player i is a pair $b_i \equiv (a_i, \pi^i_{-i})$, consisting of i 's assessment over nodes a_i , and i 's expectations of

opponents' strategies $\pi_{-i}^i = (\pi_j^i)_{j \neq i}$.⁶ We should emphasize that the assumption that player i 's expectations about his opponents' play corresponds to a strategy profile incorporates the implicit restriction that opponents randomize independently.^{7, 8}

The belief $b_i \equiv (a_i, \pi_{-i}^i)$ is *consistent* (Kreps and Wilson [17]) if $a_i = \lim_{n \rightarrow \infty} a_i^n$, where the a_i^n are obtained using Bayes rule from a sequence of strictly positive strategy profiles of the opponents, $\pi_{-i}^{i,m} \rightarrow \pi_{-i}^i$. Throughout this paper we are only concerned with consistent beliefs; at times we will refer to the sequence $\pi_{-i}^{i,m}$ as the one that makes b_i consistent.

Given a consistent belief b_i by player i , player i 's information sets give rise to a decision tree, and for each information set h_i there is a well-defined sub-tree beginning with that information set. A behavior strategy π_i is a *best response at h_i* by a player i to consistent beliefs b_i if the restriction of π_i to the subtree starting at h_i is optimal in that sub-tree. (Thus a best response at h_i supposes that the player will play optimally at subsequent nodes as well.) It is useful to define a *version* of player i , v_i , as a strategy-belief pair $v_i = (\pi_i, b_i)$. Our main solution concepts identify a *belief model*, $V = (V_1, \dots, V_n)$, that specifies a set of versions V_i for each player i .

We begin by reviewing the notion of self-confirming equilibrium and restating it in a way that is similar to our main notion.⁹

DEFINITION 2.1. Profile $\hat{\pi}$ is a *self-confirming¹⁰ equilibrium (SCE)* if there is a singleton belief model $V = \{(\pi_1, b_1)\}, \dots, \{(\pi_n, b_n)\}$ such that, for all players i

⁶ Note that what we call an "assessment" is what Kreps and Wilson [17] call a "system of beliefs for player i ," and that our "belief" is similar to what they call an "assessment." The reason we have switched terminology is that, unlike Kreps and Wilson, we consider strategic uncertainty, as reflected in the fact each player i makes his own forecast π_{-i}^i , where we do not impose $\pi_k^i = \pi_k^j$ for $i \neq j$. Thus in place of a single commonly known object (a, π) we have distinct "beliefs" $b_i \equiv (a_i, \pi_{-i}^i)$.

⁷ See Fudenberg and Kreps [10] for a discussion of this point.

⁸ Given independence, Kuhn's theorem [18] shows that there is no additional loss of generality in restricting attention to expectations that correspond to a single strategy profile π_{-i}^i , as opposed to a probability distribution over such profiles.

⁹ In Definition 2.1 every version in the belief model is required to be consistent with the overall path of play, so there is a single belief about the path of play for each player i ; this is called "unitary" self confirming in Fudenberg and Levine [12]. The alternative, "heterogeneous," version of self-confirming equilibrium only requires that (π_i, b_i) be consistent with the outcomes player i observes when playing π_i . Although heterogeneous beliefs are very important for describing some experimental outcomes, developing a "rationalizable" version of heterogeneous SCE for general games involves a number of subtleties that are beyond the scope of this paper.

¹⁰ More precisely, this is "independent unitary" self-confirming equilibrium. Since this is the primary notion we study in this paper, we omit the terms "independent unitary."

- (1) For every $(\pi_i, b_i) \in V_i$, π_i is a best response to $b_i = (a_i, \pi_{-i}^i)$ at information sets that are reached with positive probability under (π_i, π_{-i}^i) .
- (2) Every $(\pi_i, b_i) \in V_i$ has the distribution over outcomes induced by $\hat{\pi}$.

In this case we refer to $\hat{\pi}$ as a *self-confirming equilibrium* and the distribution over outcomes induced by $\hat{\pi}$ as a *self-confirming outcome*.

The other notions we develop all incorporate a *belief-closed* requirement, in the spirit of rationalizability. As in rationalizability, this requirement is intended to ensure that the strategies that player i expects player j to play could actually make sense for j to play, in the sense of being consistent with what i knows about j 's payoffs.

DEFINITION 2.2. A belief model is *belief-closed* if for every $(\pi_i, (a_i, \pi_{-i}^i)) \in V_i$, π_j^i can be generated by a mixture over strategies in the set $\{\pi_j' \mid (\pi_j', b_j) \in V_j \text{ for some belief } b_j\}$.¹¹

In words, i 's beliefs about j must be consistent with the set of j 's possible versions. Thus, the elements of the sets V_j are better viewed as "versions that player i might think player j is" than as "versions that j is likely to be." For example, if v_j 's strategy specifies an action at some off-path information set that is not optimal given j 's specified payoffs, the interpretation is that this is something i plausibly thinks that j would do if that information set is reached. As we will argue below, such beliefs can be plausible because the fact that the off-path information set was reached can lead player i to revise his beliefs about j 's payoffs.

To clarify that this condition is not driving any of the distinctions between self-confirming equilibria and our main solution concept we show that there is no loss of generality in adding the belief-closed condition to the notion of self-confirming equilibrium.

THEOREM 2.1. *Profile $\hat{\pi}$ is a self-confirming equilibrium if and only if there is a belief model V such that, for all players i ,*

- (1) *For every $(\pi_i, b_i) \in V_i$, π_i is a best response to $b_i = (a_i, \pi_{-i}^i)$ at information sets that are reached with positive probability under (π_i, π_{-i}^i) .*
- (2) *Every $(\pi_i, b_i) \in V_i$ has the distribution over outcomes induced by $\hat{\pi}$.*
- (3) *V is belief closed.*

¹¹ A behavior strategy π_j is generated by a mixture $(\alpha, 1 - \alpha)$ over π_j' and π_j'' if for every π_{-j} , the distribution over terminal nodes induced by (π_j, π_{-j}) equals the $(\alpha, 1 - \alpha)$ mixture over the distributions induced by (π_j', π_{-j}) and (π_j'', π_{-j}) , respectively.

One direction of this result is obvious: If conditions (1), (2), and (3) hold, we may pick any point (π_i, b_i) from each set V_i and Definition 2.1 will be satisfied. The converse is not quite as easy, as singleton sets satisfying (1) and (2) need not be belief closed. For example, player 1 might believe that player 2's off-path play is L, while 2's strategy specifies that 2 plays R. However, the weak optimality condition 1 does not restrict off-path play, so in this case we could add a new element to V_2 corresponding to player 1's beliefs about player 2's play. Of course, condition (1) does restrict play along the equilibrium path, but condition (2) ensures that beliefs about on-path play are correct. A formal proof along these lines is straightforward, and we omit it.

SCE imposes optimality only at information sets that are reached. Intuitively, this corresponds to the idea that the players learn the path of play but have no information about their opponents' payoffs. For example, in the game of Fig. 2.1, SCE allows 2 to play **d** so long as 2's information set is not reached in the course of play. As noted by Selten [24], player 2 can "threaten" to play **d**, and thus induce 1 to play **L**. However, this threat is not "credible" if 1 knows 2's payoff function, for then player 1 should realize that player 2 would play **u** if ever her information set is reached. For this reason, in many settings the weak rationality condition used by Nash and self-confirming equilibrium incorporates too little information about opponents' payoffs.

Thus, our first step towards introducing a theory in which players make use of information about each others payoffs is to introduce a notion of rationalizability that strengthens the optimality condition, condition 1, to require that player i 's strategy be optimal not only along the path of play, but at all information sets that are not precluded by that strategy.

DEFINITION 2.3. A belief model, V , is *rationalizable at reachable nodes* if for all i :

(1') For every $(\pi_i, b_i) \in V_i$, π_i is a best response to $b_i = (a_i, \pi_{-i}^i)$ at information sets reachable under π_i .

(3) V is belief closed.




FIGURE 2.1




FIGURE 2.2

DEFINITION 2.4. Profile $\hat{\pi}$ is a *rationalizable self-confirming equilibrium (RSCE)* if there is a belief model V that is rationalizable at reachable nodes, and such that, for all i

- (2) Every $(\pi_i, b_i) \in V_i$ has the distribution over outcomes induced by $\hat{\pi}$.

Turning back to the game in Fig. 2.1, we see that the RSCE notions capture what we wanted: **L** is not part of any beliefs that are rationalizable at reachable nodes. To see this, observe that 2's information set is always reachable, so condition (1') implies that the only strategy in V_2 is **u**. From condition (3), player 1 must believe this, and so he plays **R**. An important feature of RSCE is that a strategy need not be optimal at information sets that the strategy itself precludes. The reason that we do not wish to impose optimality at such information sets is that this stronger requirement is not robust to the presence of a small amount of payoff uncertainty. To see this, consider the game in Fig. 2.2.




FIGURE 2.3

In this game the outcome **L** occurs in the Nash equilibrium (**LD**, **d**), but not in any subgame-perfect equilibrium. However, in the game of incomplete information in Fig. 2.3, where payoffs are very likely to be as in Fig. 2.2, the outcome **L** occurs in a sequential equilibrium. So requiring optimality at all information sets rules out the outcome **L** in Fig. 2.2 but not in Fig. 2.3; hence this requirement is robust to small payoff uncertainties.¹² It is easy to see that by construction RSCE achieves our objectives in Fig. 2.2: since player 1's second information set is not reachable when 1 plays **L**, the outcome **L** can occur in a RSCE. In Section 4 we explore the relationship between RSCE and solution concepts that impose optimality at all information sets.

3. EXAMPLES

This section contains some examples that clarify the concepts defined in the preceding section.

EXAMPLE 3.1. Ordinary self-confirming equilibrium allows two players to disagree about the play of the third. This example demonstrates the intuitive idea that the possibilities for such disagreements are reduced when players must believe that opponent's play is a best response at reachable nodes. Consider the version of the extensive-form game shown in Fig. 3.1 which Fudenberg and Kreps [10] used to show that mistakes about off-path play can lead to non-Nash outcomes: In Fig. 3.1 the outcome (**A**, **a**) is self-confirming for any values of x and y . It is supported by player 1 believing that player 3 will play **R** and player 2 believing that player 3 will play **L**. However, because 3's information set is reachable, this outcome is not RSCE if both x and y have the same sign: If $x, y > 0$ then players 1 and 2 forecast that 3 will play **R**, and so 2 plays **d**; if $x, y < 0$ then 3 plays **L** so 1 plays **D**. However, if x and y have opposite signs, then (**A**, **a**) is a RSCE outcome, since players 1 and 2 are not required to have the same beliefs about player 3's off-path assessment of the relative probability of the two nodes in her information set, and player 1 can think that 3's assessment makes **R** optimal, while player 2 can think that 3's assessment induces her to play **L**.¹³

¹² Just as in previous work related to this notion of robustness, one may be able to identify a smaller set of robust predictions if one feels confident that certain forms of payoff uncertainty are much less likely than others. (We ourselves have no such confidence; we note the point because it is often raised in seminars.)

¹³ This example shows that even requiring optimality at all information sets, as in the notion of a sequentially RSCE defined in Section 4, need not be Nash.




FIGURE 3.1

EXAMPLE 3.2. The next example shows that it is possible to have outcomes that are self-confirming and rationalizable, yet fail to be RSCE. Thus the RSCE concept does more than simply take the intersection of sets that satisfy its constituent assumptions. Consider the extensive form game shown in Fig. 3.2. In this example $(\mathbf{u}, \mathbf{d}, \mathbf{U})$ is a self-confirming equilibrium (and even a Nash equilibrium) since player 2's information set is off the equilibrium path, and so he may play \mathbf{d} . Moreover the following belief model is rationalizable at reachable nodes: $V_1 = \{(\mathbf{u}, (\mathbf{d}, \mathbf{U})), (\mathbf{u}, (\mathbf{d}, \mathbf{D}))\}$, $V_2 = \{(\mathbf{a}, (\mathbf{r}, \mathbf{U})), (\mathbf{d}, (\mathbf{r}, \mathbf{D})), (\mathbf{d}, (\mathbf{u}, \mathbf{D})), (\mathbf{a}, (\mathbf{u}, \mathbf{U}))\}$, and $V_3 = \{(\mathbf{U}, (\mathbf{u}, \mathbf{a})), (\mathbf{U}, (\mathbf{u}, \mathbf{d})), (\mathbf{U}, (\mathbf{r}, \mathbf{d})), (\mathbf{D}, (\mathbf{r}, \mathbf{a}))\}$. (Note that we did not specify assessments, only strategies: for players 1 and 2 the assessments are trivial; for player 3 an assessment is needed only in the third element of V_3 . Here we specify the assessment that puts probability




FIGURE 3.2

1 on the upper node.) Thus the path (\mathbf{u}, \mathbf{U}) is possible according to each set V_i .

However, the outcome (\mathbf{u}, \mathbf{U}) is not rationalizable self-confirming. Intuitively, this is because player 1 should realize that player 2 knows that player 3 is playing up and then deduce that player 2 will play \mathbf{a} . Notice that player 2's information set is always reachable since regardless of how player 2 plays it can be reached unilaterally by player 1. Hence any beliefs for player 2 that satisfy (1') have optimization by player 2 at his information set. Moreover, the beliefs must agree with the equilibrium path, so player 2 must believe that 3 is playing \mathbf{U} . So all possible b_2 's have player 2 playing \mathbf{a} . Thus π_1 must be a best response to the belief b_1 in which 3 plays \mathbf{U} (because 1 knows the equilibrium path) and in which 2 plays \mathbf{a} (from our discussion of V_2 and the belief-closed condition), and so 1 must play \mathbf{r} instead of \mathbf{u} .

This shows that the belief-closed condition does have extra power when combined with conditions (1') (optimality at reachable nodes) and (2) (knowledge of the path), even though it is vacuous when combined with conditions (1) (optimality on the path) and (2).

EXAMPLE 3.3. We next consider further the fact that RSCE allows two players to disagree about the play of a third. Fudenberg and Levine [12] showed that in games with identified deviators the set of (unitary) SCE is not altered by adding the requirement that players have the same beliefs about one another. For RSCE, which incorporates the additional assumption of optimality at reachable nodes, this is no longer the case. For a particularly simple example consider the perfect-information game shown in Fig. 3.3.

We claim first that the outcome (\mathbf{R}, \mathbf{d}) is RSCE. (In what follows we do not specify assessments, as they are trivial.) To see this, let $V_1 = ((\mathbf{R}, \mathbf{RR}), \mathbf{d}, \mathbf{D})$, $V_2 = (\mathbf{d}, (\mathbf{R}, \mathbf{RR}), \mathbf{U})$ and $V_3 = V_1 \cup V_2$. Notice that these sets are consistent and all have the same path. Moreover condition (3) is satisfied since it only requires that if one player believes a second is playing a particular way, there is a version of the second player who is playing that way. In particular, two players may disagree about the play of a third, as is the




FIGURE 3.3

case here. It remains to check sequential optimality for each player. Player 3 is indifferent, so is certainly playing optimally. Finally, inspection shows that players 1 and 2 are both playing best responses to their beliefs at each information set.

However, this outcome cannot arise in a RSCE in which player 1 and player 2 agree about the play of player 3. With that additional requirement, an RSCE with outcome (\mathbf{R}, \mathbf{d}) would need to be subgame-perfect, since given that play there are no unreachable information sets. To see that the outcome (\mathbf{R}, \mathbf{d}) is not subgame perfect, note that any randomization by player 3 gives both player 1 or player 2 approximately the same amount, with player 2 getting slightly more than player 1. If the amount player 1 gets is strictly greater than 5 then backwards induction implies $(\mathbf{R}, \mathbf{r}, \mathbf{RR})$. If the amount player 1 gets is strictly less than 5, and player 2 gets strictly more than 0 then backwards induction implies either $(\mathbf{L}, \mathbf{r}, \mathbf{RR})$ or $(\mathbf{L}, \mathbf{r}, \mathbf{DD})$. If the amount player 1 gets is equal to 5 then we can have any mixture of $(\mathbf{R}, \mathbf{r}, \mathbf{RR})$ and $(\mathbf{L}, \mathbf{r}, \mathbf{RR})$. If the amount player 2 gets is strictly smaller than 0 backwards induction implies $(\mathbf{L}, \mathbf{r}, \mathbf{DD})$. If it is equal to zero, player 1 must play \mathbf{DD} , player 2 must play \mathbf{r} , and so player 1 plays \mathbf{L} . So (\mathbf{R}, \mathbf{d}) cannot occur in an RSCE when players 1 and 2 agree about the play of player 3.

This example relies on player 3 being indifferent, but that can be avoided by replacing 3's move with a simultaneous-move subgame between 3 and 4 that has two strict equilibria, with payoffs for 1 and 2 as in the figure. (This is a multi-stage game with observed actions and hence has identified deviators.) As in Example 3.2, RSCE allows players 1 and 2 to each expect a different Nash equilibrium in the stage game between players 3 and 4.¹⁴

EXAMPLE 3.4. If every path through the tree hits at most one information set of every player, then all information sets are reachable under any profile.¹⁵ In particular, if the game is finite and there is a unique backwards induction solution (which is true for generic assignments of payoff vectors to terminal nodes) then RSCE coincides with the backwards induction solution. This is true in particular in the game shown in Fig. 3.4.

Note that players 1 and 3 have the same payoff at every terminal node: this figure is the "agent form" of the game in Fig. 2.2, where player 1's information set 1' has been assigned to an "agent" (player 3) with the same payoffs. The reason RSCE makes different predictions in these two games is that, as we show in the next section, it captures the predictions that are

¹⁴ The same is true if RSCE is strengthened so that strategies are optimal at all information sets, as in the notion of sequentially RSCE defined in Section 4.

¹⁵ In this case rationalizability at reachable nodes coincides with its strengthening to sequential rationalizability, which requires optimality at all information sets, as defined in Section 4.




FIGURE 3.4

robust to small amounts of incomplete information *provided* that the players' doubts about their opponents' payoffs are not correlated. Thus an unexpected move by player 1 can signal that player 1's own payoffs are different than had been supposed, but does not change beliefs about the payoffs of other players. We say more about this issue of correlated payoff uncertainty below.

4. ROBUSTNESS OF RATIONALIZABLE SELF-CONFIRMING EQUILIBRIUM

Implicit in our approach is the idea that rationality at reachable nodes is more likely than at arbitrary nodes. The underlying reasoning behind this is the idea that a player's own decision to deviate does not convey to him any information about other players' rationality; while an opponent's decision to deviate may indicate a degree of irrationality. To provide a formal rationale for this reasoning, we are led to consider whether equilibrium is robust, meaning that it is not changed significantly by small perturbations in the game that is played. We do so using the model of elaborations proposed by Fudenberg *et al.* [11].

An *elaboration* with personal types of the game consists of a finite collection Θ_i of types for each player. Letting $\Theta \equiv \times_i \Theta_i$, the game tree of the elaboration has nodes (θ, x) and terminal nodes (θ, z) . There is a probability distribution μ over types Θ . Information sets in the elaboration are of the form (θ_i, h) where h_i is an information set in the original game. In particular $(\theta', x) \in (\theta_i, h_i)$ if and only if $\theta'_i = \theta_i$ and $x \in h_i$, so that each player's own type is her private information. Payoffs in the elaboration are $u_i(\theta_i, z)$, and so depend only on the terminal node of the original game and the player's type. In the elaboration, the "normal" types θ_i^0 have approximately the same payoffs as in the original game.

Let E denote the original game, and consider a sequence E^k of elaborations. We say that $E^k \rightarrow E$ if

- (1) There is a uniform bound on the number of types in Θ_i^k .
- (2) There is a uniform bound on the utility functions u_i^k .
- (3) $\lim_k u_i^k(\theta_i^0, z) = u_i(z)$.
- (4) $\lim_k \mu^k(\theta^0) = 1$.¹⁶

Suppose V_i^k are a sequence of sets of versions for player i . We say that $V_i^k \rightarrow V_i$ if

(i) for every sequence $(\pi_i^k, (a_i^k, \pi_{-i}^{ik}))$ in V_i^k and every accumulation point d of the induced sequence of distributions on terminal nodes of the original game, $D(\pi^k, \pi_{-i}^{ik})$, there exists some $(\pi_i, (a_j, \pi_{-i}^i))$ in V_i^k , with $D(\pi_i, \pi_{-i}) = d$, and

(ii) if for every $(\pi_i, (a_i, \pi_{-i}^i))$ in V_i , the induced distribution $D(\pi_i, \pi_{-i}) = d$ is the limit of some such sequence.¹⁷

Given a solution concept, we say that the solution is *robust* with respect to (independent) elaborations if whenever $E^k \rightarrow E$, and $V^k \rightarrow V$ with V^k satisfying the solution concept for the elaborated games, then V satisfies the solution concept for the original game.

THEOREM 4.1. (a) *If the belief model V is rationalizable at reachable nodes, then it is robust with respect to independent elaborations.*

(b) *If $\hat{\pi}$ is a RSCE, then it is robust with respect to independent elaborations.*

Remark 4.1. Theorem 4.1 does not extend to robustness to correlated elaborations. To see this, consider the elaboration shown in Fig. 4.1 of the game in Example 3.4 (this is essentially the same as the elaboration of Fig. 2.2 shown in Fig. 2.3).

Here the outcome \mathbf{L} can have probability close to 1 in a sequential equilibrium, and so can certainly occur in a RSCE, yet the outcome is ruled out by RSCE in the original game. Theorem 4.1 shows that RSCE is robust; we now show that it is the smallest robust concept that is at least as large as one requiring optimality at *all* information sets. The following definition strengthens the notion of rationalizability at reachable nodes to all nodes; equivalently it weakens the notion of sequential equilibrium for extensive-form games to the context of rationalizability. The subsequent definition strengthens RSCE, also by imposing optimality at all information sets.

¹⁶ Conditions (1) and (4) together imply that as k grows the event $\theta = \theta^0$ becomes common p -belief for p arbitrarily near 1 (Fudenberg and Tirole [14, Theorem 14.5]).

¹⁷ This is the notion of convergence used in Dekel and Fudenberg [9].




FIGURE 4.1

DEFINITION 4.1. The belief model V is *sequential rationalizable* if for all players i :

(1'') If $(\pi_i, b_i) \in V_i$ then π_i is a best response to $b_i = (a_i, \pi^i_{-i})$ at all information sets.

(3) V is belief closed.

DEFINITION 4.2. Profile $\hat{\pi}$ is a *sequential rationalizable self-confirming equilibrium* if there is a belief model V that is sequentially rationalizable and such that for all i :

(2) Every $(\pi_i, b_i) \in V_i$ has the distribution over outcomes induced by $\hat{\pi}$.

Greenberg [15] defines more general versions of both these concepts that do not require the game to be common knowledge. When it is, his null MACA is equivalent to sequential rationalizability and path MACA is equivalent to sequentially rationalizable self-confirming equilibrium.¹⁸

THEOREM 4.2. (a) *If V is rationalizable at reachable nodes then there exists a sequence of elaborations $E^k \rightarrow E$ and belief models $V^k \rightarrow V$ that are sequentially rationalizable in the elaborations.*

¹⁸ Greenberg does not impose common knowledge of the game because his motivation is more encompassing than ours: He “offer[s] a way to formalize and analyze social environments in which players may “live in different worlds”, but nevertheless, they often follow a “mutually acceptable course of action”—each player for his own “rational” reasons. That is, each player analyzes his own extensive form game that represents his word” [15].

(b) If $\hat{\pi}$ is RSCE, then there exists a sequence of elaborations $E^k \rightarrow E$ that have sequentially RSCE $\pi^k \rightarrow \hat{\pi}$.

Remark 4.2. This generalizes our discussion of Figures 2.2 and 2.3 concerning the non-robustness of sequentially RSCE. As in Dekel and Fudenberg [9], the proof constructs elaborations in which each player has two types, the “normal” or “sane” type and a second type that is completely indifferent between all outcomes and so use whatever “off-path” strategy that is convenient for the proof.

5. RELATED LITERATURE

Börgers [7] showed that assuming almost common knowledge of rationality and caution yields the solution concept, $S^\infty W$, that Dekel and Fudenberg [9] were led to by considerations of robustness. (This is the set left after first eliminating weakly dominated strategies, then applying iterated strict dominance.) Given the results in the preceding section it is then natural to examine and confirm the relationship between an epistemic model and RSCE. We sketch this relationship below, omitting the formal details; see Börgers and references therein for formal definitions of the concepts we use, such as Monderer and Samet’s notion [19] of almost common knowledge.

Caution means that players only use strategies that are a best reply to a full support belief. This rules out weakly dominated strategies, while RSCE, like sequential equilibrium, permits some weakly dominated strategies. Thus to obtain an equivalence with a solution concept that satisfies almost common knowledge of caution and of rationality we must strengthen RSCE as follows. A belief model, V , is *perfectly* rationalizable at reachable nodes if condition (1) in Definition 2.3 is strengthened so that not only is π_i a best response to b_i at reachable nodes, but it is also a best response to the sequence $\pi_{-i}^{i,m}$ that makes b_i consistent at those information sets. Similarly, an outcome $\hat{\pi}$ is a *perfectly* RSCE if in addition it satisfies condition (2), that the distribution induced by the versions $(\pi_i, b_i) \in V_i$ have the distribution over outcomes induced by $\hat{\pi}$.¹⁹

The only remaining difference between strategies in a belief model that is perfectly rationalizable at reachable nodes and those in $S^\infty W$ arises

¹⁹ One might wonder whether the additional requirement of perfection is robust with respect to elaborations. It is, if we weaken the notion of elaboration by strengthening condition 3 in the definition above of convergence of an elaboration to require that $u_i^k(\theta_i^0, z) = u_i(z)$ instead of $\lim_k u_i^k(\theta_i^0, z) = u_i(z)$. Without strengthening the notion of convergence, the impact of the trembles in breaking ties can be undone by small perturbations in the payoffs. (See Dekel and Fudenberg [9].)

because Börgers allows for correlation while we do not. In particular, the two coincide for two-person games. The effect of correlation is twofold. In Example 3.4, a game of perfect information where each player moves only once, rationalizability at reachable nodes rules out the outcome **L**, but this outcome survives $S^\infty W$, since even after **D** is deleted for player 3 by weak dominance, 2's choice of **d** is not strictly dominated. Intuitively, this reflects the fact that **L** can be justified by an elaboration with correlated types, so that a deviation by player 1 could convince player 2 that player 3's payoff function is different than had been originally supposed. In addition to the possibility of correlated perturbations, $S^\infty W$ also allows players to believe that their opponents' play is correlated. It is well known that allowing for this larger set of beliefs results in different and larger solution sets.²⁰ Since this and other effects of correlation are well understood, we have chosen not to develop them formally here. The characterization of perfectly RSCE results from adding the requirement that the distribution over outcomes is almost common knowledge.

Rubinstein and Wolinsky [23] define a related solution concept, rationalizable conjectural equilibrium, or RCE, for games in strategic form. The main distinction between RCE and this paper is our focus on the extensive form: Our model therefore restricts behavior at (some) off-path information sets, which theirs does not. In addition they allow for correlation, while we assume independence. Finally, the papers use different formulations of the idea that beliefs must accord with observed play: Where we suppose that players observe terminal nodes, they allow observations to be generated by more general "signal functions"; on the other hand, we allow players to observe distributions, while they consider only deterministic observations. To best see the relationship between their work and ours consider two-person games, to set aside the difference due to correlation, and restrict attention to the common case where a deterministic path is observed. In this case their solution concept is the same as Battigalli and Guatoli's conjectural equilibrium (CE) [2] and self-confirming equilibrium, both of which assume only rationality rather than common certainty of rationality. Our focus in this paper was to *add* to self-confirming equilibrium (robust) elements of common knowledge of extensive-form rationality, which obviously are not contained at all in Rubinstein and Wolinsky's RCE, and *a fortiori* in Battigalli and Guatoli's CE.

²⁰ For example, consider the three-player game in Fudenberg and Kreps [10], where player 1 has a choice of either playing "out" and ending the game or playing a simultaneous-move subgame with players 2 and 3. If "out" is not a best response to any strategy profile of 2 and 3, yet "out" is a best response to a correlated strategy, then "out" cannot be played in any self-confirming equilibrium with independent beliefs, but it can be played if correlated beliefs are permitted. Fudenberg and Kreps also discuss the interpretation of correlated beliefs in the context of learning in games.

TABLE 1

| | Rationality | (Almost) ^a Common certainty of rationality | ACC of rationality and caution |
|---------------------------------|--|---|---|
| No observation | S | $S^\infty{}^b$ | $S^\infty W$ = perfectly rationalizable at reachable nodes ^c |
| Paths are publicly observed. | Self confirming equilibria = CE ^d = rationalizable CE ^e | | Perfect RSCE |

^a The parentheses around almost are meant to indicate that in this column the result does not depend on whether or not almost is included; the column on the right requires the restriction to almost common certainty since caution is introduced.

^b See Bernheim [4], Pearce [20], and Tan and Werlang [25].

^c See Börgers [7].

^d See Battigalli and Guatoli [2].

^e See Rubinstein and Wolinsky [23].

To summarize we present Table 1 relating the solution concepts discussed in this section for the case of two-person games when deterministic paths of play in the extensive-form are observable.

APPENDIX

THEOREM 4.1. (a) *If the belief model V is rationalizable at reachable nodes, then it is robust with respect to independent elaborations.*

(b) *If $\hat{\pi}$ is a RSCE, then it is robust with respect to independent elaborations.*

Proof. (a) Suppose that $E^k \rightarrow E$, and $V^k \rightarrow V$ where V^k is rationalizable at reachable nodes. We will show that the V 's are consistent and satisfy (1''). Condition (3) is obviously preserved in the limit.

First observe that the consistency of V_i is immediate, since we can take a diagonalization of the sequences of full-support sequences corresponding to the V_i^k 's. To prove (1''), suppose that $(\pi_i, (a_i, \pi_{-i}^i)) \in V_i$ and π_i is not a best response to $(\pi_i, (a_i, \pi_{-i}^i))$ at some h_i that is reachable under π_i . Since $V^k \rightarrow V$, π_i is a limit point of a sequence $\{\pi^k\}$ with $(\pi_i, (a_i, \pi_{-i}^i)) \in V_i^k$; moreover if h_i is reachable under π_i , then it is reachable under π^k . We claim that the conditional probability $pr^k(\theta_i^0 | h_i, \pi^k)$ that i is "normal" at such an h_i must converge to 1 in any sequence of consistent assessments.

Because the game has perfect recall, if h_i is not reached under π^k but is reachable for player i , there is no strategy for player i such that h_i is reached under π^k_{-i} . Consequently $pr^k(\theta_i^0 | h_i, \pi^k) = \mu^k(\theta_i^0) \rightarrow 1$. But then the fact that $\pi^k_i | h_i$ is optimal in the subtree that describes the game from h_i on given beliefs b_i^k implies that $\pi_i | h_i$ is optimal with respect to b_i following h_i , so that the property of being rationalizable at reachable nodes is robust.

(b) Fix a profile $\hat{\pi}$ that is a RSCE or a heterogeneous RSCE, and suppose that $E^k \rightarrow E$. Let $V^k \rightarrow V$ be sets that satisfy (1''), (2), and (3). Part (a) shows that (1'') and (3) are preserved in the limit, and it is obvious that (2) is as well. ■

THEOREM 4.2. (a) *If V is rationalizable at reachable nodes then there exists a sequence of elaborations $E^k \rightarrow E$ and belief models $V^k \rightarrow V$ that are sequentially rationalizable in the elaborations.*

(b) *If $\hat{\pi}$ is RSCE, then there exists a sequence of elaborations $E^k \rightarrow E$ that have sequentially RSCE $\pi^k \rightarrow \hat{\pi}$.*

Proof. (a) Let V be rationalizable at reachable nodes. We will consider elaborations E^k in which each player has one alternative type who is completely indifferent over outcomes. We construct V_i^k in two steps. First, we modify $v_i \equiv (\pi_i, (a_i, \pi^i_{-i})) \in V_i$ so that it is sequentially rational at all information sets in the elaborations V_i^k : the rational type's play is modified at all unreachable information sets so as to be sequentially rational given $b_i = (a_i, \pi^i_{-i})$; the indifferent type plays π . This gives rise to a strategy $\tilde{\pi}_i^k(v_i)$ in the elaborated game.

We now construct the V_i^k . Given $v_i \equiv (\pi, (a_i, \pi^i_{-i})) \in V_i$ and j we can find (due to condition (3)) a probability distribution $p_j(\cdot | v_i)$ over V_j consistent with π^i_{-i} . We may also combine the behavior strategies $\tilde{\pi}_j^k(v_j)$ using the weights $p_j(\cdot | v_i)$ to get behavior strategies $\bar{\pi}_j^{ik}$ for k th elaboration. The set V_i^k will consist of all (a_i^k, π^k) that are generated by some $v_i = (\pi_i, (a_i, \pi^i_{-i}))$ in V_i , where (i) $\pi_j^k = \bar{\pi}_j^{ik}$ for $j \neq i$, (ii) $\pi_i^k = \tilde{\pi}_i^k(v_i)$, and (iii) a_i^k is defined as follows. Let π^n_{-i} be the full support beliefs that makes the assessment a_i consistent. Let $\pi^{k,n}_{-i}$ be a strategy profile for i 's opponents in the elaboration given by the sane opponents playing their sequentially rational strategies (the $\tilde{\pi}_j^k(v_j)$) and the indifferent types playing π^n_{-i} . Finally, note that while $\pi^{k,n}_{-i}$ does not have full support, it is the case that against $\pi^{k,n}_{-i}$ every information set of i can be reached by some strategy of i , and therefore (under perfect recall) $\pi^{k,n}_{-i}$ uniquely determines an assessment for i .

That the sets V_i^k are sequentially rationalizable follows immediately from their construction.

(b) Suppose now that $\hat{\pi}$ is RSCE. Then there is a belief model V that is rationalizable at reachable nodes and also satisfies condition (2) or (2').

Construct the elaborations and the sets $V^k \rightarrow V$ as above. Since the same type's play along the path of $\hat{\pi}$ is the same in the elaborations as in the original game, and the indifferent type's play in each $\tilde{\pi}_i^k(v_i)$ is the same as in the strategy π_i that generated it, the path of play is $\hat{\pi}$ in each of the elaborations. In other words conditions (2) and (2') are inherited by the beliefs constructed in the elaborations. ■

REFERENCES

1. K. Basu, Strategic irrationality in extensive games, mimeo Princeton University, 1985.
2. P. Battigalli and D. Guatoli, Conjectural equilibria and rationalizability in a macroeconomic game with incomplete information, extended abstract, University Bocconi, 1988.
3. E. Ben-Porath, Rationality, Nash equilibrium and backwards induction in perfect information games, *Rev. Econ. Stud.* **64** (1997), 23–46.
4. D. Bernheim, Rationalizable strategic behavior, *Econometrica* **52** (1984), 1007–1028.
5. G. Binmore, Modeling rational players, I and II, *Econ. Philos.* **3** and **4** (1987–1988), 179–214 and 9–55.
6. L. Blume and W. R. Zame, The algebraic geometry of perfect and sequential equilibrium, *Econometrica* **62** (1988), 783–794.
7. T. Börgers, Weak dominance and approximate common knowledge, *J. Econ. Theory* **64** (1994), 265–276.
8. M. Costa-Gomez, V. Crawford, and B. Bruseta, Cognition and behavior in normalform games: An experimental study, mimeo <http://weber.ucds.edu/Depts/Econ/Wpapers/Files/ucsd9822.pdf>, 1995.
9. E. Dekel and D. Fudenberg, Rational behavior with payoff uncertainty, *J. Econ. Theory* **52** (1990), 243–267.
10. D. Fudenberg and D. M. Kreps, Learning in extensive games. I. Self-confirming equilibrium, *Games Econ. Behav.* **8** (1995), 20–55.
11. D. Fudenberg, D. M. Kreps, and D. K. Levine, On the robustness of equilibrium refinements, *J. Econ. Theory* **44** (1988), 354–380.
12. D. Fudenberg and D. K. Levine, Self-confirming equilibrium, *Econometrica* **61** (1993), 523–546.
13. D. Fudenberg and D. K. Levine, Steady state learning and Nash equilibrium, *Econometrica* **61** (1993), 547–574.
14. D. Fudenberg and J. Tirole, “Game Theory,” MIT Press, Cambridge, MA, 1991.
15. J. Greenberg, Towering over Babel: Worlds apart-but acting together, mimeo McGill University, 1996.
16. F. Gul, Rationality and coherent theories of strategic behavior, *J. Econ. Theory* **70** (1996), 1–31.
17. D. Kreps and R. Wilson, Sequential equilibria, *Econometrica* **50** (1982), 863–894.
18. H. Kuhn, Extensive games and the problem of information, in “Annals of Mathematics Studies,” Vol. 28, Princeton Univ. Press, Princeton, NJ, 1953.
19. D. Monderer and D. Samet, Approximating common knowledge with common beliefs, *Games Econ. Behav.* **1** (1989), 170–190.
20. D. Pearce, Rationalizable strategic behavior and the problem of perfection, *Econometrica* **52** (1984), 1029–1050.
21. P. Reny, Rationality in extensive-form games, *J. Econ. Perspectives* **6** (1992), 103–118.

22. P. Reny, Common belief and the theory of games with perfect information, *J. Econ. Theory* **59** (1993), 257–274.
23. A. Rubinstein and A. Wolinsky, Rationalizable conjectural equilibrium: Between Nash and rationalizability, *Games Econ. Behav.* **6** (1994), 299–311.
24. R. Selten, Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit, *Z. Ges. Staatswiss.* **121** (1965), 301–324.
25. T. C. C. Tan and S. R. C. Werlang, The Bayesian foundations of solution concepts of games, *J. Econ. Theory* **45** (1988), 370–391.