

BACKWARD INDUCTION IN GAMES WITHOUT PERFECT RECALL

JOHN HILLAS AND DMITRIY KVASOV

ABSTRACT. The equilibrium concepts that we now think of as various forms of backwards induction, namely subgame perfect equilibrium (Selten, 1965), perfect equilibrium (Selten, 1975), sequential equilibrium (Kreps and Wilson, 1982), and quasi-perfect equilibrium (van Damme, 1984), are explicitly restricted to games with perfect recall. In spite of this the concepts are well defined even in games without perfect recall. There is now a small literature examining the behaviour of these concepts in games without perfect recall.

We argue that in games without perfect recall the original definitions are inappropriate. Our reading of the original papers is that the authors were aware that their definitions did not require the assumption of perfect recall but they were also aware that without the assumption of perfect recall the definitions they gave were not the “correct” ones. We give definitions of two of these concepts, sequential equilibrium and quasi-perfect equilibrium, that identify the same equilibria in games with perfect recall and behave well in linear games without perfect recall. We also extend these definitions to nonlinear games. Finally we give the appropriate redefinition of perfect equilibrium in games without perfect recall for both linear and nonlinear games.

1. INTRODUCTION

The game theorists who defined the equilibrium concepts that we now think of as various forms of backwards induction, namely subgame perfect equilibrium (Selten, 1965), perfect equilibrium (Selten, 1975), sequential equilibrium (Kreps and Wilson, 1982), and quasi-perfect equilibrium (van Damme, 1984), explicitly restricted their analysis to games with perfect recall. In spite of this the concepts are well defined, exactly as they defined them, even in games without perfect recall. There is now a small literature examining the behaviour of these concepts in games without perfect recall. Jeff Kline (2005) looks at what happens in games without perfect recall to solutions defined in exactly the same way as they were defined in games with perfect recall. Joe Halpern and Rafael Pass (2017) modify the definitions of van Damme and Kreps and Wilson in a somewhat different manner than we do. Adrian Marple and Yoav Shoham (2013) also define related concepts, though their approach involves splitting the owners of information sets, and even the owners of some nodes within an information set into separate players.

We shall argue that in games without perfect recall the original definitions are inappropriate. Our reading of the original papers is not that the authors were unaware that their definitions did not require the assumption of perfect recall, but rather that they were aware that without the assumption of perfect recall the definitions they gave were not the “correct” ones. In this paper we give definitions of these concepts, that identify the same equilibria in games with perfect recall and behave well in games without perfect recall. By “behave well” we mean exhibit the same inclusions as the original concepts exhibit in games with perfect recall,

Date: IRNLP Version 4 May 27, 2018.

Key words and phrases. extensive form games; perfect recall; sequential equilibrium; quasi-perfect equilibrium.

namely, perfect and quasi-perfect equilibria should be sequential equilibria which, in turn, should be Nash equilibria. Moreover in generic games perfect, quasi-perfect, and sequential equilibria should coincide.

We shall, below, first define the concepts of sequential equilibria and quasi-perfect equilibria. We focus on these concepts because they seem to us to be the right concept for backward induction and a combination of backward induction and admissibility respectively. Later we redefine the concept of extensive form perfect equilibrium for games without perfect recall.

In the next section we shall briefly develop the notation we use and review some of the facts about the equilibria of the various classes of games.

2. EXTENSIVE FORM GAMES: NOTATION AND BASIC RESULTS

We shall be using the model of extensive form games developed by Kuhn (1950, 1953) who modified and generalised the definition given by von Neumann and Morgenstern (1947). Kuhn distinguished between games with perfect recall, in which players remember at each occasion they move everything they knew and did in the past, and games where this is not true. An implication of a player having perfect recall is that any play of the game will cut each of the player's information sets at most once. Kuhn made this requirement part of his definition of an extensive form game. Isbell (1957) expanded the class of games by dropping the requirement that each play of a game cut each information set at most once. Isbell called games in which this requirement was met *linear games* and the more general class of games where this condition was not assumed *nonlinear games*. This more general definition was later, under the name "repetitive games," considered by Alpern (1988), and more famously, under the name "absent-mindedness," by Piccione and Rubinstein (1997a), and, following them, by many others, (Gilboa, 1997; Battigalli, 1997; Grove and Halpern, 1997; Halpern, 1997; Lipman, 1997; Aumann, Hart, and Perry, 1997a,b; Piccione and Rubinstein, 1997b). In this paper we shall use Isbell's terminology and distinguish between linear and nonlinear games. We thus have three successively broader classes of extensive form games: games with perfect recall, linear games, and nonlinear games.

A definition of an extensive form game starts with the notion of a game tree. There are two equivalent ways of defining a game tree. The first, used by Kuhn, defined a game tree as a finite connected graph with no loops and a distinguished initial node or root. More recently it has become typical to define a game tree as a partially ordered finite set of nodes or equivalently in terms of a *predecessor* function, which specifies the node that comes immediately before a given node. These ways of defining a game tree are equivalent and given one definition one can define the elements of the other definition. We shall use the terminology from each definition as convenient. We let the finite set of nodes be X , the initial node x_0 and the (immediate) *predecessor* function

$$p : X \rightarrow X \cup \{\emptyset\},$$

the function that for every node gives the node that comes immediately before it in the game tree (with $p(x_0) = \emptyset$ and x_0 is the only node for which this is true). When thinking of the game tree as a graph the set $\{p(x), x\}$ is a branch. We say that t is a terminal node if there is no node $x \in X$ with $p(x) = t$. We partition X into the sets T of terminal nodes and D of nonterminal or decision nodes.

As well as the game tree an extensive form game Γ consists of: a finite set of players, $N = \{1, 2, \dots, N\}$; a set of actions, A , and a labeling function $\alpha : X \setminus \{x_0\} \rightarrow A$ where $\alpha(x)$ is the action at $p(x)$ that leads to x such that if $p(x) = p(x')$ and $x \neq x'$ then $\alpha(x) \neq \alpha(x')$; a collection of information sets, \mathcal{H} , and a function $h : D \rightarrow \mathcal{H}$

that assigns for every decision node which information set the node is in; a function $n : \mathcal{H} \rightarrow \hat{N} = N \cup \{0\}$, where player 0 is Nature and $n(h)$ is the player who moves at information set h and $\mathcal{H}_n = \{h \in \mathcal{H} \mid n(h) = n\}$, the information sets controlled by player n ; a function $\rho : \mathcal{H}_0 \times A \rightarrow [0, 1]$ giving the probability that action a is taken at the information set $h \in \mathcal{H}_0$; and functions (u_1, \dots, u_N) with $u_n : T \rightarrow \mathbb{R}$ being the payoff to player n . We shall also consider the collection of non-empty subsets of \mathcal{H}_n which we shall denote $\bar{\mathcal{H}}_n$. An element H of $\bar{\mathcal{H}}_n$ is some collection of information sets of Player n .

We let $A(x) = \{a \in A \mid a = \alpha(x') \text{ for some } x' \text{ with } p(x') = x\}$. That is, $A(x)$ is the set of choices that are available at node x . If x is a terminal node then $A(x) = \emptyset$. If $h(x) = h(x')$, then $A(x) = A(x')$. For convenience and notational simplicity we assume that if $h(x) \in \mathcal{H}_0$ and $h(x) = h(x')$ then $x = x'$, that is that all $h \in \mathcal{H}_0$ are singletons.

Definition 1. If no path from x_0 to a terminal node cuts any information set more than once then we call the game a *linear game*.

The idea of perfect recall is that a player has perfect recall if, at each of his information sets he remembers what he knew and what he did in the past. This idea was introduced and formally defined by Kuhn (1950, 1953). Later Selten (1975) gave an equivalent definition that is perhaps closer to the intuitive idea. Ritzberger (1999) and, even more fundamentally, Alós-Ferrer and Ritzberger (2016, 2017) redefined extensive form games and gave a deep discussion of the meaning of perfect recall. For our purposes the standard formulation of extensive form games and definition of perfect recall are sufficient. We give here the definition of perfect recall by Selten.

Definition 2. A player is said to have *perfect recall* if whenever that player has an information set containing nodes x and y and there is a node x' of that player that precedes node x there is also a node y' in the same information set as x' that precedes node y and the action of the player at y' on the path to y is the same as the action of the player at x' on the path to x .

We now define the various notions of strategy that we shall use.

Definition 3. A *pure strategy* in an extensive form game for Player n is a function that maps each of his information sets to one of the actions available at that information set. We denote the set of Player n 's pure strategies by S_n , the set of *pure strategy profiles* by $S = \times_{n \in N} S_n$, and the set of *extended pure strategy profiles* by $\hat{S} = \times_{n \in \hat{N}} S_n$.

Definition 4. A *behaviour strategy* in an extensive form game for Player n is a function that maps each of his information sets to a probability distribution over the actions available at that information set. We denote the set of Player n 's behaviour strategies by B_n and the set of *behaviour strategy profiles* by $B = \times_{n \in N} B_n$.

Definition 5. A *mixed strategy* in an extensive form game for Player n is a probability distribution over the player's pure strategies. We denote the set of Player n 's mixed strategies by Σ_n and the set of *mixed strategy profiles* by $\Sigma = \times_{n \in N} \Sigma_n$.

In nonlinear games we also need to consider randomisations over behaviour strategies. Different terms have been used to refer to such strategies in the literature. Isbell called them mixed strategies and called what we, and most of the rest of the literature, call mixed strategies linear strategies; Alpern called them randomised strategies; Selten called them behaviour strategy mixtures; and others have called them mixtures of behaviour strategies. We follow Mertens, Sorin, and

Zamir (2015) in calling them general strategies. We *can* consider such strategies for linear games, but we do not need to do so.

Definition 6. A *general strategy* in an extensive form game for Player n is a probability distribution over the player's behaviour strategies. We denote the set of Player n 's general strategies by G_n and the set of general strategy profiles by $G = \times_{n \in N} G_n$.

In linear decision problems, that is, one player linear games the player always has a pure strategy that is at least as good as any more general strategy. In nonlinear decision problems this is not the case. A player may be able to do better with a behaviour strategy than he can with any pure strategy. In games with more than one player equilibria may involve players randomising, or, equivalently, players being uncertain of what other players are choosing. In linear games this means that we look for equilibria in mixed strategies, in nonlinear games we look for equilibria in general strategies. Kuhn showed that in games with perfect recall one could achieve the same uncertainty about the other players with behaviour strategies as one could with mixed strategies and so, in such games one could instead look for equilibria in behaviour strategies. We can be a little more explicit. We first define the notion of the Kuhn-equivalence of two strategies.

Definition 7. Given two strategies of Player n , x and y in $S_n \cup B_n \cup \Sigma_n \cup G_n$ we say that x is *Kuhn-equivalent* to y if for any general strategy profile g_{-n} in G_{-n} of the other players the profiles (x, g_{-n}) and (y, g_{-n}) induce the same distribution on the terminal nodes.

And we can now formally state Kuhn's Theorem.

Kuhn's Theorem. *If Player n has perfect recall then for any mixed strategy σ_n in Σ_n there is a behaviour strategy b_n in B_n that is Kuhn-equivalent to σ_n . In a linear game for any player n in N and for any behaviour strategy b_n in B_n there is a mixed strategy σ_n in Σ_n that is Kuhn-equivalent to b_n .*

The first part of Kuhn's Theorem was stated and proved by Kuhn (1953). The second part is almost implicit in Kuhn's paper and was formally stated and proved by Isbell (1957).

Remark 1. When we defined perfect recall we defined what it meant for a player to have perfect recall and Kuhn's Theorem refers to the equivalence of behaviour and mixed strategies for any player who has perfect recall. When we define linear games we do not do so player by player and Kuhn's Theorem only refers to linear games in which all players satisfy the requirement that their information sets are cut at most once by a path through the game. This is consistent with the way things are normally done in the literature. It is however true that if a player satisfies the requirement that none of her information sets are cut more than once by any path through the tree then for any behaviour strategy of that player the player has a Kuhn-equivalent mixed strategy, whether or not the other players satisfy this requirement. It might, perhaps, be better to define "linear players" and state Kuhn's Theorem in this way and indeed Mertens, Sorin, and Zamir (2015) present the results this way.

Since there is an infinite number of behaviour strategies the space of general strategies is infinite dimensional. Fortunately, we do not need to consider all general strategies. The following result allows us to restrict ourselves to a finite dimensional subset of G_n . This result was proved by Alpern (1988).

Proposition 1 (Alpern 1988). *For any player n in N there is a finite number K_n such that for any general strategy g_n of Player n there is general strategy g'_n putting probability on only K_n elements of B_n that is Kuhn-equivalent to g_n .*

Proof. Each terminal node t in T defines a set of decision nodes of Player n on the path from the initial node to t . For each of these nodes there is a branch $\{x, y\}$ from x with y also on the path to t . A behaviour strategy b_n of Player n induces a conditional probability on the branch $\{x, y\}$ conditional on x having been reached. Let $q_n(t, b_n)$ be the product of the conditional probabilities generated by b_n on the branches following nodes owned by Player n that occur on the path to t . If Player n has no nodes on the path to t we let $q_n(t, b_n) = 1$. Similarly define $q_0(t, b_0)$ for Nature, where b_0 is Nature's only strategy. Thus if the players play $b = (b_1, b_2, \dots, b_N)$ the probability that terminal node t will be reached is $\prod_{n \in \hat{N}} q_n(t, b_n)$. Notice that b_n influences the distribution over terminal nodes only through $q_n(t, b_n)$.

Let

$$Q_n = \{(q_t)_{t \in T} \subset [0, 1]^T \mid \text{for some } g_n \text{ in } G_n \text{ for all } t \ q_t = \int_{B_n} q_n(t, b_n) dg_n(b_n)\}.$$

It is clear that Q_n is the convex hull of those points $(q_t)_{t \in T}$ in Q_n with g_n putting probability only on one behaviour strategy, that is, of the set

$$C_n = \{(q_t)_{t \in T} \subset [0, 1]^T \mid \text{for some } b_n \text{ in } B_n \text{ for all } t \ q_t = q_n(t, b_n)\}.$$

But since C_n (and Q_n) are subsets of \mathbb{R}^T , by Carathéodory's Theorem any q in Q_n can be written is a convex combination of at most $T + 1$ elements of C_n . That is it is generated by a general strategy g_n that puts probability on at most $T + 1$ elements of B_n .

We have shown that for any g_n we can find a g'_n that puts probability on only $T + 1$ elements of B_n such that g_n and g'_n generate the same element of Q_n . But g_n will influence the probability of a final node only through the element of Q_n it generates and the result follows. \square

Remark 2. In our proof we have given $T + 1$ as the bound on the number of behaviour strategies that may receive positive probability. This can be substantially strengthened. In general many different terminal nodes may be associated with the same set of edges following nodes of Player n on the path to that terminal node. We would need in Q_n only one dimension for each such set of edges.




FIGURE 1. The absent-minded driver (Piccione and Rubinstein 1977).

Remark 3. It is not true that we can restrict attention to only a fixed finite subset of B_n . Consider the Absent-minded Driver example in Figure 1. We claim that there is no finite set of behaviour strategies such that the outcome from any behaviour strategy can be replicated by some mixture over the given set. Suppose that we have T behaviour strategies b^1, b^2, \dots, b^T with $b^t = (x^t, 1 - x^t)$ with x^t being the probability that the player chooses L . So, if the player plays b^t he ends up with outcome a with probability x^t , with outcome b with probability $x^t(1 - x^t)$, and with outcome c with probability $(1 - x^t)^2$. Let \bar{x} be the smallest value of x^t strictly greater than 0.

Consider the behaviour strategy $b^0 = (\bar{x}/2, 1 - (\bar{x}/2))$. This strategy gives outcome a with probability $\bar{x}/2$ and outcome b with probability $(\bar{x}/2)(1 - (\bar{x}/2))$. Now for any b^t which gives strictly positive probability of outcome b we have that the ratio of the probability of outcome a to the probability of outcome b is $1/(1 - x^t) \geq 1/(1 - \bar{x})$.

Thus if we have a general strategy putting probability only on b^1, b^2, \dots, b^T that gives outcome b with the same probability as b^0 , that is with probability $(\bar{x}/2)(1 - (\bar{x}/2))$ it will give outcome a with probability at least

$$\left(\frac{1}{1 - \bar{x}}\right) \left(\frac{\bar{x}}{2}\right) \left(1 - \frac{\bar{x}}{2}\right) = \left(\frac{\bar{x}}{2}\right) \left(\frac{2 - \bar{x}}{2 - 2\bar{x}}\right) > \frac{\bar{x}}{2},$$

and so it does not induce the same probabilities on outcomes as b^0 .

As a consequence of Proposition 1, instead of working with the infinite dimensional space G_n we can instead work with the finite dimensional space

$$\hat{G}_n = \Delta_{K_n} \times B_n^{K_n},$$

the Cartesian product of the K_n -simplex with K_n copies of B_n . The typical element $(\alpha_1, \dots, \alpha_k, \dots, \alpha_{K_n}, b_n^1, \dots, b_n^k, \dots, b_n^{K_n}) \in \hat{G}_n$ means that for each k Player n plays his behaviour strategy b_n^k with probability α_k . For every element of G_n there is a Kuhn-equivalent element in the subset \hat{G}_n .

And this allows us to avoid certain technical issues and, more importantly, to use techniques of real algebraic geometry to prove the generic equivalence of sequential and quasi-perfect equilibria.

We also need to define the completely mixed general strategies.

Definition 8. A general strategy g_n in G_n of Player n is *completely mixed* if, for any open subset O of B_n , $g_n(O) > 0$. We denote the set of all completely mixed general strategies of Player n by G_n^0 and the set of completely mixed profiles by G^0 .

We also define the corresponding subset of \hat{G}_n .

Definition 9. A general strategy g_n in \hat{G}_n of Player n is *completely mixed* if there is some g'_n in G_n^0 such that g'_n is Kuhn-equivalent to g_n . We denote the set of all such strategies of Player n by \hat{G}_n^0 and the corresponding set of profiles by \hat{G}^0 .

In the next section we look at a number of examples that give a good indication of some of the necessary features of the definitions of quasi-perfect and sequential equilibria in games without perfect recall.

3. NECESSARY FEATURES OF THE DEFINITION

In this section we shall look at a number of examples that illustrate some of the issues that arise in defining the concepts that we are interested in in games without perfect recall. Most of the central issues arise already in linear games and all of the example we consider in this section are linear games.

We first look at the game given in extensive form in Figure 2 in normal form in Figure 3. This is a very slight modification of a game considered by Kuhn (1953). This is a two player zero-sum game in which each player has a unique optimal strategy and so there is a unique equilibrium in mixed strategies. Player 1's optimal strategy is not equivalent to any behaviour strategy.




FIGURE 2. A game without perfect recall.

		Player 2	
		<i>IN</i>	<i>OUT</i>
Player 1	<i>In, S</i>	0, 0	$\frac{3}{2}, -\frac{3}{2}$
	<i>In, D</i>	0, 0	$-\frac{1}{2}, \frac{1}{2}$
	<i>Out, S</i>	$-\frac{3}{2}, \frac{3}{2}$	0, 0
	<i>Out, D</i>	$\frac{1}{2}, -\frac{1}{2}$	0, 0

FIGURE 3. The corresponding normal form game.

Consider the normal form of the game given in Figure 3. Note that for Player 1, the strategies (In, D) and (Out, S) are strictly dominated. Once the dominated strategies are removed the game is similar to matching pennies and the unique equilibrium is

$$\left\{ \left(\frac{1}{4}, 0, 0, \frac{3}{4} \right), \left(\frac{3}{4}, \frac{1}{4} \right) \right\},$$

a mixed strategy profile in which Player 1 is playing a strategy that is not equivalent to any behaviour strategy.

Among other things this implies that there is no equilibrium in behaviour strategies. For, if there were then the equivalent mixed strategies would be such that neither player had a behaviour strategy that he preferred—this is the definition of an equilibrium in behaviour strategies—and thus no pure strategy that he preferred. Thus that profile of mixed strategies would be an equilibrium, contradicting the fact that this game has a unique equilibrium in mixed strategies in which one of the players plays a strategy that is not equivalent to any behaviour strategy.

Let us think for a moment why there is, in this example, no equilibrium in behaviour strategies. In the mixed equilibrium Player 1 at his first information set sometimes plays *In* and sometimes plays *Out* and at his second information set sometimes plays *S* and sometimes plays *D*. However he coordinates his choices so that if he plays *In* at his first information set he plays *S* at his second information set and if he plays *Out* at his first information set he plays *D* at his second information set. It is not possible to achieve such coordination using behaviour strategies.

Another candidate for an equilibrium in behaviour strategies is the profile of behaviour strategies corresponding to an equilibrium of the agent normal form. The equilibria of the agent normal form are (In, S, IN) , (Out, D, OUT) and $\{(x, 1 - x), (1/4, 3/4), (x, 1 - x) \mid 0 \leq x \leq 1\}$. Why are these (or at least the equivalent behaviour strategies) not equilibria? Player 1 could deviate to a different behaviour strategy that involves different behaviour at *both* of his information sets.

We should not be surprised at the nonexistence of equilibria in behaviour strategies. The definition of such equilibria allows coordination by the player in his deviations that he is not permitted in his equilibrium strategies.

If we want a solution where such coordination is not permitted then the appropriate solution is equilibria of the agent normal form, and the perfect equilibria—since each player will have only one information set there is no need to distinguish between extensive form perfect, normal form perfect, or quasi perfect—of the game will encompass whatever aspects of backward induction we want. In this paper we shall consider solutions in which such coordination among the choices of a player at his different information sets is possible.

In the game we have been considering there is a unique equilibrium and in that equilibrium all information sets are reached with positive probability. There is no need for backward induction arguments. Nevertheless it is instructive to look at the example to see what the nature of a backward induction requirement will be. In the example all information sets are reached so there should be no issue about what the beliefs of the players will be.

Lets consider the beliefs of Player 1 at his second information set. It's clear that his beliefs will differ depending on what pure strategy he is playing. If he is playing (In, S) then he will assess a probability $3/7$ on the left node and $4/7$ on the right node, while if he is playing the strategy (Out, D) then he will assess a probability 1 on the left node and 0 on the right node. Moreover his assessment of what the other player is playing is also different. When Player 1 is playing the strategy (Out, D) then if called upon to move at his second information set he will assess a probability of 1 on Player 2 having played *IN* rather than his prior probability of $3/4$.

In games with perfect recall there is an equivalence, roughly corresponding to the equivalence between mixed and behaviour strategies shown by Kuhn, between beliefs on the nodes of the information set and beliefs on what the other players are playing. In games without perfect recall this is not so. It will be clear as we consider other examples that beliefs about which node of an information set a player is at are not rich enough to encompass what is needed in games without perfect recall. Thus we shall think of a Player's assessment as a belief about the strategies of the other players, including Nature. Here, when Player 1 is playing (In, S) , his beliefs will be that with probability $4/7$ Nature chose "Right" (and Player 2 chose *IN* with probability $3/4$ and *OUT* with probability $1/4$) and with probability $3/7$ Nature chose "Left" and Player 2 choose *IN*. On the other hand, if he is playing (Out, D) then he will assess a probability 1 on the fact that Nature chose "Left" and Player 2 chose *IN*.

Let us summarise the general features illustrated by the analysis in the previous paragraph of our particular example. A player's beliefs at an information set may depend on which pure strategy he himself is playing. And, for each of his pure strategies, his beliefs about the strategies of the others may not be an independent product (across players) of probability distributions on the pure strategies of the other players. Of course, this latter fact is also true in games with perfect recall.




FIGURE 4. No One-Deviation Principle.

Before proceeding further with our consideration of how we should define beliefs we need to consider another aspect in which the situation differs between games with perfect recall and games without perfect recall. In games with perfect recall in order to show that a strategy is optimal, against beliefs consistent with it, it is enough to show that at each information set it is optimal taking as fixed the choice at the other information sets of the player (see Hendon, Jacobsen, and Sloth, 1996; Perea, 2002). In games without perfect recall this is not true.

Consider the game of Figure 4. Player 1 decides whether to take an outside option or allow the two agents of Player 2 to play a coordination game in which Player 1 obtain the same payoff as Player 2.

Consider the profile in which Player 1 chooses T and Player 2 chooses DR . This is an equilibrium. Moreover, if we consider only deviations at one information set at a time, it appears to robustly satisfy backward induction type arguments. Nevertheless, in the one person subgame beginning with Player 2's first move there is only one equilibrium, namely UL , and hence the only subgame perfect equilibrium is (B, UL) .

Thus there seems to be no hope that we can satisfy the “one deviation principle” that several backward induction concepts satisfy in games with perfect recall.

Nevertheless when we consider a player at a particular information set we cannot allow that player to deviate at any arbitrary collection of other information sets following that one, or, at least, that such a requirement is problematic. In the game of Figure 5 consider the equilibrium (T, RU) . We claim that this is a reasonable equilibrium, and in particular that it intuitively satisfies a reasonable notion of backward induction. However, if we allow Player 2 at the node at which he chooses between L and R to deviate at both information sets that would upset this equilibrium.




FIGURE 5. Cannot always deviate at multiple information sets.

One option would be to allow players to deviate at an information set and at information sets all of whose nodes follow some node in the information set under consideration. Such a requirement is however not strong enough. Consider again the game of Figure 2. Since neither information set of Player 1 follows the other in this sense we cannot at either information set have Player 1 consider deviating at both of his information sets. Thus we would allow a strategy profile equivalent to the equilibrium of the agent normal form, which we have argued we do not wish to do. One could perhaps get around this by adding dummy moves at various points throughout the game. We take a somewhat different approach, which seems to us a little more attractive. When we consider deviations we will allow the player to deviate at an arbitrary subset of her information sets, but require her to update conditional on the event that one of the information sets in that subset of information sets had been reached. Thus in the equilibrium (T, RU) considered above Player 2 can deviate at both her information sets, but when she does her beliefs about what Player 1 is playing should remain the Player 1 is playing T since that is consistent with the fact that one of Player 2's information sets was reached.

In the standard definition of beliefs, such as in Kreps and Wilson (1982), the beliefs of a player, conditional on an information set being reached, are given by a probability distribution over the nodes of that information set. We have, for the most part, in this section referred instead to beliefs about what the other players are playing. The game shown in Figure 6 shows that we need to do so.

All equilibria of this game involve Player 1 playing X . What should Player 2 believe at her information set? Notice that simply having some distribution over the nodes of her information set does not tell Player 2 what to do. She also needs to know what Player 1 will do at his second information set. Now, if we gave Player 1 a behaviour strategy we would specify a behaviour at Player 1's second information set. But this is not what Player 2's beliefs should be. The pure strategies TD and BU are strictly dominated by TU and BD respectively, and both are strictly dominated by X . Thus Player 2 may be uncertain whether she is at the left node or at the right node, but should believe that if she is at the left node then Player 1 will play U at his second information set and that if she is at the right node then Player 1 will play D at his second information set. An efficient way of describing these beliefs would be that at her information set Player 2 believes that Player 1




FIGURE 6. Need Beliefs About Strategies

played either TU or BD . (A little further analysis shows that her beliefs must be that these are equally likely.)

Thus we see that a player’s beliefs at an information set, or a collection of information sets since, as we have seen earlier, we must allow a player to deviate at a collection of information sets, must encompass not just where the player is in the information set, but also what strategies the others are playing. When we are working with mixed strategies and the players beliefs are about the pure strategies of the others the players’ beliefs about the strategies of the others will be concentrated on strategies that reach the information sets in question and so generate the distribution over the nodes of the information sets.

We shall, in Section 5, provide definitions appropriate for nonlinear games. However, most of the issues that arise already arise in linear games and, since that setting is both more familiar and simpler, we shall develop our definitions there first.

4. DEFINITIONS AND RESULTS FOR LINEAR GAMES

We shall now define sequential equilibria and quasi-perfect equilibria. Since we have seen that we cannot hope to satisfy a one-deviation property and that it will be necessary to consider players deviating simultaneously at a number of information sets we shall define beliefs not at an information set but at a collection of information sets. In the original definition of sequential equilibrium beliefs were defined as a probability distribution over the nodes of an information set. Here we define beliefs as distributions over the pure strategies that are being played, including Nature’s strategy.

Definition 10. A *system of beliefs* μ defines, for each n in N and each H in $\bar{\mathcal{H}}_n$, a distribution $\mu(s_0, s_1, \dots, s_N | H)$ over the extended profiles of pure strategies that reach H . Given μ we also consider $\mu_{S_n}(s_n | H)$ the marginal distribution on S_n given H and $\mu_{S_{-n}}(s_{-n} | s_n, H)$ the conditional distribution on S_{-n} conditional on s_n and H .

Recall that we have seen above that a player’s beliefs at an information set about what strategies the other players are playing may differ depending on what pure strategy he himself is playing.

We first define sequential equilibria.

Definition 11. A pair (σ, μ) is a *consistent assessment* supported by a sequence of completely mixed strategy profiles σ^t in Σ if $\sigma^t \rightarrow \sigma$ with μ^t a system of beliefs obtained from σ^t as conditional probabilities and $\mu^t \rightarrow \mu$.

Definition 12. A pair (σ, μ) is *sequentially rational* if, for each n , for each H in $\bar{\mathcal{H}}_n$, and for each s_n in S_n , if $\mu_{S_n}(s_n | H) > 0$ then s_n maximises

$$E_{\mu_{S_n}(s_{-n}|s_n, H)} u_n(t_n, s_{-n})$$

over the set of all t_n in S_n such that t_n differs from s_n only at information sets in H .

Definition 13. A pair (σ, μ) is a *sequential equilibrium* if it is both consistent and sequentially rational.

Quasi-perfect equilibria are defined in a similar way.

Definition 14. A strategy profile σ is a *quasi-perfect equilibrium* if there is a sequence of completely mixed strategy profiles $\sigma^t \rightarrow \sigma$ with (σ, μ) a consistent assessment supported by σ^t and (σ, μ^t) sequentially rational for all t .

Observe that the definitions of sequential equilibrium and quasi-perfect equilibrium differ only in requiring that (σ, μ^t) being sequentially rational for all t rather than only (σ, μ) being sequentially rational.

We now give a number of results about the concepts we have defined. The first two results say that in games with perfect recall we obtain the “same” equilibria as the original definitions.

Proposition 2. *If the game has perfect recall then for any sequential equilibrium (σ, μ) there is a behaviour strategy profile b , Kuhn-equivalent to σ , that is the strategy profile part of a sequential equilibrium according to the definition of Kreps and Wilson (1982). Moreover for any sequential equilibrium in the sense of Kreps and Wilson there is a Kuhn-equivalent mixed strategy profile σ and a system of beliefs μ such that (σ, μ) is a sequential equilibrium.*

Proof. We first show that if (σ, μ) is a sequential equilibrium then there is $(b, \tilde{\mu})$ with b Kuhn-equivalent to σ that is a Kreps-Wilson sequential equilibrium.

Suppose that the game has perfect recall and that (σ, μ) is a sequential equilibrium. Thus there is σ^t , completely mixed, with $\sigma^t \rightarrow \sigma$, μ^t obtained from σ^t as conditional probability, $\mu^t \rightarrow \mu$, and (σ, μ) sequentially rational in the sense of Definition 12.

Now, by Kuhn’s Theorem, there are behaviour strategies b^t Kuhn-equivalent to σ^t and b Kuhn-equivalent to σ . Let $\tilde{\mu}^t$ be the Kreps-Wilson beliefs obtained from b^t as conditional probabilities. Taking subsequences if necessary, we have $\tilde{\mu}^t \rightarrow \tilde{\mu}$, and, moreover, $\tilde{\mu}$ will give the same distribution over the nodes of each information set H as would be implicit in $\mu(\cdot | H)$. (And so, in games with perfect recall for any sequential equilibrium (σ, μ) , $\mu_{S_n}(\cdot | s_n, H)$ is independent of s_n for all s_n that make H possible.)

Now $(b, \tilde{\mu})$ is, by construction, Kreps-Wilson consistent. And, if it were not Kreps-Wilson sequentially rational then there would be some information set h such that b put positive probability on some action a that did not maximise against $(b, \tilde{\mu})$. But then it must be that there is some pure strategy s_n that takes action a at h with $\mu_{S_n}(s_n | \{h\}) > 0$. But then there would be some pure strategy that differed from s_n only at h that did better against $\mu_{S_n}(\cdot | s_n, \{h\})$ than s_n , contradicting that (σ, μ) is a sequential equilibrium.

Now suppose that $(b, \tilde{\mu})$ is a Kreps-Wilson sequential equilibrium supported by the sequence $(b^t, \tilde{\mu}^t)$. Let σ^t be completely mixed and Kuhn-equivalent to b^t and

μ^t be obtained from σ^t as conditional probabilities. Again taking subsequences if necessary, let $\sigma^t \rightarrow \sigma$ and $\mu^t \rightarrow \mu$. We claim that (σ, μ) is a sequential equilibrium.

For any player n and for any $H \in \bar{\mathcal{H}}_n$ the beliefs $\mu(\cdot | H)$ generated by σ^t will be the same as those that would have been generated by b^t , since σ^t and b^t are Kuhn-equivalent. Moreover, since the game has perfect recall and the beliefs are consistent with the strategy profile, the one-deviation principle holds. Thus we need consider only single information sets $\{h\}$ with $h \in \mathcal{H}_n$.

Consider an information set $h \in \mathcal{H}_n$. We claim that for any s_n such that $\mu_{S_n}(s_n | \{h\}) > 0$ it cannot be that there is a better t_n that differs from s_n only at h . For if there were then the action taken at h by s_n must be both taken with positive probability at h by b_n and worse than some other action at h against $(b, \tilde{\mu})$, contradicting that $(b, \tilde{\mu})$ is a Kreps-Wilson sequential equilibrium. Thus (σ, μ) is a sequential equilibrium, as required. \square

Proposition 3. *If the game has perfect recall then for any quasi-perfect equilibrium σ there is a behaviour strategy profile b , Kuhn-equivalent to σ , that is a quasi-perfect equilibrium according to the definition of van Damme (1984). Moreover, for any quasi-perfect equilibrium in the sense of van Damme there is a Kuhn-equivalent mixed strategy profile σ that is a quasi-perfect equilibrium.*

Proof. The proof is almost exactly the same as the proof of Proposition 2. \square

The next two results say that the relation between the concepts is as it was in games with perfect recall.

Proposition 4. *Any quasi-perfect equilibrium is a sequential equilibrium strategy profile.*

Proof. This result follows from the observation following Definition 14 and the fact that $E_{\mu_{S_{-n}}^t(s_{-n}|s_n, H)} u_n(t_n, s_{-n})$ is continuous in $\mu_{S_{-n}}^t(s_{-n} | s_n, H)$. \square

Proposition 5. *For any extensive game form, except for a semialgebraic set of payoffs (to the terminal nodes) of lower dimension than the set of all payoffs, every sequential equilibrium is a quasi-perfect equilibrium.*

Proof. The proof follows in a straightforward way similar to the proof of the generic equivalence of perfect and sequential equilibria in Blume and Zame (1994) and the proof, based on Blume and Zame, of the generic equivalence of quasi-perfect and sequential equilibria in Hillas, Kao, and Schiff (2017) or Pimienta and Shen (2013). \square

Finally we have the result proved for games with perfect recall by van Damme (1984) and Kohlberg and Mertens (1986) relating quasi-perfect and sequential equilibria to proper equilibria (Myerson, 1978) of the normal form.

Proposition 6. *Every proper equilibrium is a quasi-perfect equilibrium (and hence a sequential equilibrium).*

Proof. Suppose that σ is a proper equilibrium supported by the sequence σ^t of completely mixed strategies with σ^t an ε^t -proper equilibrium. Let μ^t be obtained from σ^t as conditional probability and, taking subsequences, if necessary, let $\mu^t \rightarrow \mu$.

Let $H \in \bar{\mathcal{H}}_n$ and s_n and t_n such that they differ only at information sets in H . Thus if t_n is better than s_n against $\mu_{S_{-n}}^t(\cdot | s_n, H)$ then, since t_n differs only at information sets in H , it must be that t_n is better than s_n against σ^t and, from the definition of ε -proper equilibria, $\sigma_n^t(s_n) < \varepsilon^t \sigma_n^t(t_n)$ and so $\mu_{S_n}^t(s_n | H) \rightarrow 0$ or $\mu_{S_n}(s_n | H) = 0$. Thus σ is a quasi-perfect equilibrium, as required. \square

Since every game has a proper equilibrium this result also implies the existence of sequential and quasi-perfect equilibria.

5. EXTENSION TO NONLINEAR GAMES

We now give the definitions of sequential and quasi-perfect equilibria appropriate for nonlinear games. If we look, as we would in defining consistent assessments, at the limit of beliefs defined as the limit of beliefs derived as conditional probability from completely mixed general strategies we may obtain, at some information sets, beliefs that do not put positive probability on behaviour strategies that reach the information sets under consideration with positive probability.




FIGURE 7. In a nonlinear game a consistent assessment may not reach H .

Consider the nonlinear game given in Figure 7. Suppose we take a sequence of general strategies g^t such that g^t puts all probability on the single behaviour strategy profile $b^t = ((p^t, 1 - p^t), (q^t, 1 - q^t))$ with $(p^t, q^t) \rightarrow (0, 0)$. In the limit this is equivalent to the pure strategy profile (B, R) . This clearly should not be a sequential equilibrium—Player 2's behaviour clearly should not be viewed as sequentially rational against B . And yet, since the strategy B does not reach Player 2's information set, R is a best response against B .

Thus we need to add, to the beliefs over the strategies, beliefs over the information sets under consideration and, when we define sequential rationality, make our utility comparisons conditional on the collection of information sets. When we could rely on the strategies given positive probability in the beliefs reaching the information sets under consideration with positive probability, the beliefs over the information sets were implicit in the beliefs over the strategies. When the strategies may not reach the information sets this is no longer the case.

In defining quasi-perfect equilibria we require best responses against the beliefs μ^t which put probability only on strategies that do reach the information sets on which we are conditioning. However the strategies we test for optimality are given by the limit belief μ and, while this does not matter in the example of Figure 7, it

is not hard to find an example in which it will matter. Consider the game given in Figure 8.




FIGURE 8. A consistent assessment where a player’s own strategy may not reach H .

Consider a sequence g^t where g^t puts probability 1 on the single profile of behaviour strategies $b^t = ((b_{11}^t, b_{12}^t), b_2^t) = (((\delta^t, 1 - \delta^t), (\delta^t, 1 - \delta^t)), (\delta^t, 1 - \delta^t))$ where b_{11} gives the distribution over $\{T, B\}$ at Player 1’s first information set, b_{12} gives the distribution over $\{U, D\}$ at Player 1’s second information set, and b_2 gives the distribution over $\{L, R\}$ at Player 2’s information set, and $\delta^t \rightarrow 0$. Clearly the beliefs conditional on Player 1’s second information set put probability 1 on $g^t \rightarrow g$ where g puts probability 1 on (a behaviour strategy profile Kuhn-equivalent to) the pure strategy profile (BD, R) . But BD does not reach Player 1’s second information set and so if we tried to define quasi-perfect equilibrium as we did for linear games there would be no strategy that does better than BD against R . Yet (BD, R) does not seem to be a backward induction profile. At Player 1’s second information set U is clearly better than D . We need specify a distribution over the nodes of the information set—in this case, since there is only one node in the information set, that is quite trivial—and test whether the expected payoff conditional on this collection of information sets being reached is better than that obtained by some other behaviour strategy that differs only at nodes in this collection of information sets.

In linear games and when considering a single information set the beliefs over the information set are straightforward to generate from a strategy profile that reaches the information set; a strategy profile gives a distribution over the terminal nodes and, in a linear game, each terminal node is associated with at most one node of an information set. Thus we obtain a conditional distribution over the nodes of the information set. In nonlinear games or when considering more than one information

set the definition is just a little less straightforward. We use a modification of an idea introduced by Grove and Halpern (1997) and Halpern (1997). They define the upper frontier of an information set as the set of nodes that do not have any node in the same information set preceding it. We extend this idea to collections of information sets.

Definition 15. Given a collection of information sets $H \in \bar{\mathcal{H}}_n$ the *initial frontier* of H is

$$\hat{H} = \{x \in H \mid \text{there is no } x' \text{ in } H \text{ that precedes } x\}.$$

Notice that the initial frontier of H is not the union of the upper frontiers of the information sets in H . If there are more than one information set in H then a node in one information set may be preceded by a node in another information set in H . Thus \hat{H} may differ from H even in linear games, and indeed even in games with perfect recall, though in that case we do not need to consider collections of information sets.

We now define the beliefs. We first amend the definition of beliefs we gave for linear games to the case of nonlinear games. This essentially involves replacing pure strategies by behaviour strategies and mixed strategies by general strategies. We then extend the definition of beliefs by adding a distribution over the frontier of the collection of information sets. Finally we describe how we calculate the expected utility.

Definition 16. A *system of beliefs* μ defines, for each n in N and each H in $\bar{\mathcal{H}}_n$ a finite distribution $\mu(s_0, b_1, \dots, b_N \mid H)$ over the choice of Nature and the profiles of behaviour strategies. Given μ we also consider $\mu_{B_n}(b_n \mid H)$ and $\mu_{B_{-n}}(b_{-n} \mid b_n, H)$ the marginal distribution on B_n given H and the conditional distribution on B_{-n} conditional on b_n and H .

It is convenient to work not just with finite distributions but with distributions that can be embedded in some fixed finite dimensional space, just as we defined the space of general strategy profiles above. We need a bit more than general strategies since, once we condition on a particular information set, or collection of information sets, the distribution over the players behaviour strategies may not be independent. However, again in a similar way to the way we defined the restriction of the set of general strategy profiles, \hat{G} , we can define the restricted set of beliefs

$$\hat{M} = \Delta_{\prod_{n \in N} K_n} \times \left(\prod_{n \in N} B_n^{K_n} \right),$$

where $K_0 = |S_0|$, the cardinality of the pure strategy set of Nature. An element $\hat{\mu} \in \hat{M}$ specifies for each player n a selection of K_n behaviour strategies and a distribution over the profiles $(s_0, b_1, b_2, \dots, b_N)$ where, for each n the strategy b_n varies over the K_n specified behaviour strategies, and s_0 varies over S_0 .

A player's beliefs at an information set about what strategies the other players are playing may differ depending on what behaviour strategy he himself is playing. Notice also that, just as we did for linear games, we include the (pure) strategy of Nature in the list of strategies over which Player n has beliefs.

Definition 17. An *extended system of beliefs* $(\mu, \tilde{\mu})$ is a pair where μ is a system of beliefs as defined in Definition 16 and, for each n in N , each H in $\bar{\mathcal{H}}_n$, and each b_n in B_n such that $\mu_{B_n}(b_n \mid H) > 0$, $\tilde{\mu}$ gives a distribution $\tilde{\mu}(\cdot \mid b_n, H)$ over \hat{H} .

Definition 18. A pair $(g, (\mu, \tilde{\mu}))$ is a *consistent assessment* supported by a sequence of completely mixed general strategies g^t in \hat{G}^0 if $g^t \rightarrow g$ with $(\mu^t, \tilde{\mu}^t)$ a system of beliefs obtained from g^t as conditional probabilities and $(\mu^t, \tilde{\mu}^t) \rightarrow (\mu, \tilde{\mu})$.

Remark 4. Because $g^t \in \hat{G}$ and μ^t is derived from g^t as conditional probability we will have $\mu^t \in \hat{M}$ and so for any sequence g^t there will be some subsequence such that both g^t and μ^t converge, since both \hat{G} and \hat{M} are compact.

We are now in a position to define the expected payoff conditional on reaching a particular collection of information sets H and sequential rationality would be that at no collection of information sets of Player n could the player gain by deviating at those information sets. To do this formally we require a little more notation. Our definitions follow the definition of Kreps and Wilson (1982), appropriately modified to apply to nonlinear games.

For any node x in the game tree we let $p_1(x) = p(x)$, the immediate predecessor of x and inductively define $p_k(x) = p_1(p_{k-1}(x))$; $T(x)$ the set of terminal nodes that follow x ; $T(H)$ the set of terminal nodes that follow some node in H .

The expected utility of a player deviating at a collection of information sets H will depend on both the particular behaviour strategy she is deviating from, which will influence her beliefs, and the behaviour strategy to which she deviates which will determine her actions at the information sets in H . Thus we define the probability over terminal nodes conditional on a collection of information sets H of Player n with beliefs $(\mu, \tilde{\mu})$ who deviates from behaviour strategy b_n to β_n . If the terminal node t is not in $T(H)$ then $P^{\mu, \tilde{\mu}, b_n, \beta_n}(t | H) = 0$. If t is in $T(H)$ then there is a unique value of k such that $p_k(t)$ is in \hat{H} and we have

$$P^{\mu, \tilde{\mu}, b_n, \beta_n}(t | H) = \tilde{\mu}(p_k(t)) E_{\mu(b_{-n}|b_n, H)} \prod_{i=1}^k \beta_{n(p_i(t))}(\alpha(p_i(t)), h(p_i(t))),$$

where

$$\beta_m(\alpha, h) = \begin{cases} \beta_n(\alpha, h), & \text{if } m = n \\ b_m(\alpha, h), & \text{otherwise.} \end{cases}$$

Having defined $P^{\mu, \tilde{\mu}, b_n, \beta_n}(\cdot | H)$ we also define the conditional expectation with respect to this distribution and denote it $E^{\mu, \tilde{\mu}, b_n, \beta_n}(\cdot | H)$.

Definition 19. A pair $(g, (\mu, \tilde{\mu}))$ is *sequentially rational* if, for each n , for each H in \mathcal{H}_n , and for each b_n in B_n if $\mu_{B_n}(b_n | H) > 0$ then b_n maximises

$$E^{\mu, \tilde{\mu}, b_n, \beta_n}(u_n(z) | H)$$

over the set of all β_n in B_n such that β_n differs from b_n only at information sets in H .

Definition 20. A pair $(g, (\mu, \tilde{\mu}))$ is a *sequential equilibrium* if it is both consistent and sequentially rational.

Definition 21. A strategy profile g is a *quasi-perfect equilibrium* if there is a sequence of completely mixed general strategy profiles g^t with $(g, (\mu, \tilde{\mu}))$ a consistent assessment supported by g^t , with $(\mu^t, \tilde{\mu}^t)$ obtained from g^t as conditional probabilities, and $(g, (\mu^t, \tilde{\mu}^t))$ sequentially rational for all t .

Just as in Section 4 we showed that quasi-perfect and sequential equilibria as defined in that section were, in games with perfect recall, equivalent to the standard solutions we now wish to show that in a linear game the definitions we have now given will define solutions that are equivalent to the solutions defined in Section 4.

One does need to take a little care in defining the beliefs in the Section 4 definitions. First there is no unique map from behaviour strategies to mixed strategies and so no unique map from beliefs over behaviour strategies to beliefs over pure strategies. Kuhn (1953, page 211) does however give a canonical map that maps a behaviour strategy to one of the mixed strategies that is Kuhn-equivalent to it.

So given a belief μ over behaviour strategies we can define an equivalent belief $\bar{\mu}$ over pure strategies, though there may be many other different beliefs over pure strategies that are also equivalent.

The second issue is that we cannot go directly from the limit belief μ over behaviour strategies to the required equivalent belief over pure strategies. As we have seen, the definition we use for nonlinear games allows the limit belief to put probability only on behaviour strategies that do not reach the information sets on which we are conditioning. Thus when we define the Section 4 beliefs we first generate beliefs over S from the completely mixed beliefs μ^t . We then condition on the set of pure strategies that reach the collection H . We call this distribution $\bar{\mu}^t$. Then, taking subsequences if needed, we let $\bar{\mu}$ be a limit of this sequence.

Proposition 7. *In a linear game, if g is a quasi-perfect equilibrium then any mixed strategy profile σ that is Kuhn-equivalent to g is a quasi-perfect equilibrium as defined in Definition 14.*

Proof. □

Proposition 8. *In a linear game, if $(g, (\mu, \tilde{\mu}))$ is a sequential equilibrium, supported as a consistent assessment by the sequences g^t and $(\mu^t, \tilde{\mu}^t)$, σ is a mixed strategy profile that is Kuhn-equivalent to g and $\bar{\mu}$ a distribution on S defined from the sequence μ^t as described above then $(\sigma, \bar{\mu})$ is a sequential equilibrium as defined in Definitions 11, 12, and 13.*

Proof. □

Proposition 9. *In any finite game there exists a quasi-perfect equilibrium.*

Proof. □

Proposition 10. *Any quasi-perfect equilibrium is a sequential equilibrium strategy profile.*

Proof. The continuity of $E^{\mu, \tilde{\mu}, b_n, \beta_n}(u_n(z) | H)$ in $\mu, \tilde{\mu}$ implies that if $(g, (\mu^t, \tilde{\mu}^t))$ is sequentially rational for all t then $(g, (\mu, \tilde{\mu}))$ sequentially rational, and the result follows. □

6. PERFECT EQUILIBRIUM

In this section we give a definition of perfect equilibria. In the case of perfect equilibria there are no issues that are better illuminated in linear games and we just give a definition that applies to both linear and nonlinear games. As we shall see, this definition is more straightforward than the redefinitions of sequential and quasi-perfect equilibria. We start by defining a perturbed game. We do this in such a way that when the players use behaviour strategies it results in essentially the same perturbations as used by Selten, and others following him. However we shall define perturbed games in such a way that a perturbed game is a finite extensive form game. We do so by adding the “mistakes” explicitly as moves of Nature rather than as restrictions on some strategy space.

For a given extensive form game Γ a perturbation, δ , gives, for each information set h in \mathcal{H}_n and each action a in $A(h)$ available at h , a positive number δ_{ha} such that for each h in \mathcal{H}_n we have $\sum_{a \in A(h)} \delta_{ha} < 1$. A δ -perturbed game starts with a move of Nature. We index the moves of Nature at this initial node by γ . If D is the set of decision nodes of Γ and $A(d)$ is the set of actions available at d then the index γ specifies for each d in D a value γ_d being an element of $\{x\} \cup A(d)$. Thus there are $\prod_{d \in D} (|A(d)| + 1)$ different values of γ .

Each γ indexes a copy of the original game except that in the copy at each decision node d the action taken at d is the action taken by the player who moves

at d if $\gamma_d = x$ and the action γ_d otherwise. The probability that Nature chooses γ at the initial node is given by $\rho(\gamma) = \prod_{d \in D} \delta_{h(d)\gamma_d}$ where if γ_d is in $A(d)$ then $\delta_{h(d)\gamma_d}$ is already defined and if $\gamma_d = x$ then $\delta_{h(d)x}$ is defined as $1 - \sum_{a \in A(d)} \delta_{h(d)a}$. An information set in the perturbed game is the union over all the copies of the original tree of the nodes corresponding to the nodes in a particular information set of the original game. Informally this means that we assume that no player “sees” anything directly about the chosen γ , though, since γ does affect the actions “taken” by the players they may see some aspects indirectly. This is completely consistent with the way perturbed games are traditionally treated though one may wonder if we should allow a player to “see” whether, in the past he played an action because he chose to or the action was played in spite of the fact that he chose to take some other action. We argue below, in the context of a simple example, that this difference does not really matter.

One might wonder why it is necessary to have a possible version of the game for each decision node rather than just for each information set. Indeed for linear games we could have the possibility of a mistake determined just once for each information set. However for nonlinear games if one node in an information set follows another in the same information set the determination of whether a mistake occurs must be determined independently for each node. The definition we have given does no harm (beyond making the perturbed game rather complicated) when nodes in the same information set do not follow each other and is necessary when they do.

We now illustrate the definition for a simple example. Consider the game given in Figure 9. In the perturbed game there will be 27 copies of the game indexed $xxx, Txx, Bxx, xUx, \dots, BDR$. We illustrate part of the perturbed game in Figure 10. We now consider the issue we alluded to earlier. Consider the second information set of Player 1 containing the nodes u, v , and w . The node v follows the choice T at Player 1’s first information set while u and w follow the choice B . Thus, in the perturbed game, Player 1 does not have perfect recall. Moreover one might argue that it would be more accurate to split this information set into two, one for the nodes following the (attempted) choice of T by Player 1 and another for the nodes following the choice B . Doing this would also lead to perturbations of games with perfect recall also having perfect recall. However there are a number of disadvantages in defining perturbed games in this way. The strategy spaces of the perturbed games would now be different to the strategy spaces of the unperturbed game. They would also be more complicated to describe and the analysis would be more intricate. It is also inconsistent with the implicit treatment of the issue in Selten’s original definition of a perturbed game, and indeed in the treatment of all those who followed him in defining concepts that depended on perturbations that restricted the set of behaviour strategies available. In these definitions a player saw only the realisation of the behaviour strategy and not the action whose probability the chosen strategy maximised, that is, the strategy that the player “tried” to play.

We also claim that the argument that it is more natural to put the node v into a different information set to that containing u and w depends on a particular interpretation of the original unperturbed game. The game of Figure 9 tells us that when Player 1 moves the second time and chooses between U and D she understands exactly her situation, but it does not tell us why. One natural interpretation is that she remembers that she moved before and chose B and observed Player 2 choosing R . Another interpretation is that she neither remembers moving before nor seeing Player 2 move but understands the game and knows that if she is being asked to choose between U and D it must be that she earlier chose B and that Player 2




FIGURE 9. A Simple Game




FIGURE 10. A Perturbation of the Simple Game

then chose R . With this interpretation one would naturally place u , v , and w in the same information set.

We also argue that it does not matter which way we define the perturbed game. Consider the perturbed game as we have defined it with all of the second nodes of Player 1 in the same information set. If we give Player 1 at her second node an additional piece of information, namely whether or not she tried to play T at her first node we are giving her a piece of information that no other player will observe and that does not affect the distribution over the outcomes a , b , c , and d . While she could certainly now make a different choice depending on what she had tried to do at her first node that would have exactly the same effect on the final distribution over terminal nodes as making a random choice at the one information set and we already allow her to make random choices.

It is now straightforward to define perfect equilibria. If the original game was a linear game then the perturbed game will be a linear game. If the original game

was a nonlinear game then the perturbed game will be a nonlinear game. Moreover the information sets will be the same in the original game and the perturbed game, as will the actions available at those information sets. Thus the strategy sets will be the same in the original game and the perturbed game. And the set \hat{G} will not depend on δ .

Definition 22. We say that $g \in \hat{G}$ is a *perfect equilibrium* if there is a sequence of perturbations $\delta^t \rightarrow 0$ with $\delta^t > 0$ and $g^t \in \hat{G}$ an equilibrium of the δ^t -perturbed game with $g^t \rightarrow g$.

Remark 5. If the original game was linear then we could work with the mixed strategy space Σ instead of \hat{G} . And, if the original game had perfect recall then a perfect equilibrium would be Kuhn-equivalent to a behaviour strategy that would be a perfect equilibrium according to Selten's definition. Since \hat{G} is finite dimensional the existence of perfect equilibria follows in the standard way.

We would obtain exactly the same perfect equilibria if instead of perturbing the game by explicitly adding the "mistakes" to the extensive form we restricted the space of behaviour strategies for each player by requiring that each action be played with some small probability as is more typical in definitions of perfect equilibrium. Halpern and Pass (2017, page 14) give a definition of perfect equilibrium this way.

One issue that arises in defining perfect equilibrium in nonlinear games that does not arise in linear games has to do with the fact that in a nonlinear game a player may want to play exactly a particular completely mixed behaviour strategy. Defining perturbations in the way in which we have in this section allows a player to do so. It is not obvious to us that this is actually consistent with the idea that players make small mistakes. An alternative approach would be to require that the player play a completely mixed general strategy, that is, a strategy in G_n^0 , as we defined earlier. We shall not formally define such a concept but we shall argue that doing things in this way would actually define a different solution.

Consider the Absent-Minded Driver game given in Figure 1 with the player being Player 1 and Player 1's utility being 1 at outcome b and 0 everywhere else. Player 1 clearly maximises his expected utility by maximising the probability of b and he does this by taking actions L and R each with probability one half. This leads to outcome b with probability $1/4$ and this is the most that the probability of b can be. Now suppose that a second player, Player 2, plays after Player 1, without seeing anything that Player 1 did, choosing between A and B , with Player 2's choice not affecting Player 1's utility. If the outcome of Player 1's choices was b then Player 2 obtains 0 from A and 3 from B ; otherwise he obtains 1 from A and 0 from B . If the probability of outcome b is strictly less than $1/4$ then Player 2's only best response is to play A . If the probability of b is exactly $1/4$ then any choice of Player 2 is a best response.

So, if we define perfect equilibrium as we did in Definition 22 then the perfect equilibria are $\{(1/2, 1/2), (x, 1 - x) \mid 0 \leq x \leq 1\}$ while if we made a definition requiring that the perturbed general strategy profile be in G^0 then the total probability of outcome b would be strictly less than $1/4$ from any perturbed strategy and the only perfect equilibrium would be $((1/2, 1/2), (1, 0))$.

7. GENERAL FORM PERFECT EQUILIBRIUM

An important equilibrium refinement in linear games is that of normal form perfect equilibrium which might be thought of as a relatively strong form of admissibility. In the *definition* of an extensive form game strategies are not defined and payoffs or utilities are associated to terminal nodes. To define the normal form we define strategies and derive the utility of profiles of strategies from the utilities of

the terminal nodes. For the games they defined, a subset of the games that we are calling linear games, von Neumann and Morgenstern (1947) argued that what they called the normalised form, which we usually now call the normal form, is sufficient for analysis. Their argument applies equally well to the class of linear extensive form games. However it does not apply to nonlinear games, which von Neuman and Morgenstern did not consider.

In an extensive form game each profile of pure strategies generates a distribution over the terminal nodes of the game tree and so to each profile of pure strategies we can associate an expected payoff. Thus the normal form of a game consists of the triple (N, S, u) where N is the set of players, inherited directly from the extensive form game, $S = \times_{n \in N} S_n$ the set of pure strategy profiles, and $u : S \rightarrow \mathbb{R}^N$ the utility function giving, for each profile of pure strategies the expected utility of each player.

We extend the utility function to profiles of mixed strategies by taking expectations. An equilibrium of a normal form game consists of a profile of mixed strategies such that for each player, given the profile of mixed strategies of the others his mixed strategy is at least as good as any other of his mixed strategies. In a linear game, by Kuhn's Theorem, for any behaviour strategy there is a mixed strategy that, for any profile of strategies of the others, generates the same distribution over the terminal nodes and hence the same expected payoff. Thus if there is no mixed strategy that does better than the player's equilibrium strategy then there is not any behaviour strategy that does better. Moreover neither is there any general strategy that does better. Given a general strategy, that is a probability distribution over the set of behaviour strategies one could replace each behaviour strategy with a Kuhn-equivalent mixed strategy and take the same distribution over those mixed strategies which would, by construction, give the same distribution over terminal nodes as the original general strategy. But a probability distribution over the probability distributions over the pure strategies is equivalent, in terms of the eventual probability of each pure strategy, as some single direct distribution over the pure strategies, that is a mixed strategy. Thus, in a linear game, if there is no mixed strategy that does better than the equilibrium strategy, neither is there any general strategy that does better. In a linear game an equilibrium in mixed strategies is not vulnerable to deviations to general strategies.

The same is not true in nonlinear games. One could still define the normal form in the same way and indeed still define equilibria in mixed strategies. However it would no longer be true that equilibria in mixed strategies were not vulnerable to deviations to behaviour strategies. And so the normal form is not a good representation of a nonlinear game and an equilibrium in mixed strategies is not a suitable solution concept. Instead we define what we call the general form of the game.

Definition 23. The *general form* of an extensive form game is a triple (N, B, u) where N is the set of players, $B = \times B_n$ is the set of behaviour strategy profiles and $u : B \rightarrow \mathbb{R}^N$ the utility function giving, for each profile of behaviour strategies the expected utility of each player.

Again, we extend the utility function to profiles of general strategies by taking expectations. An equilibrium of a general form game consists of a profile of general strategies such that for each player, given the profile of general strategies of the others his general strategy is at least as good as any other of his general strategies. Just as in a normal form game if no pure strategy does better than a mixed strategy then no mixed strategy will either, in a general form game if no behaviour strategy does better than a general strategy then no general strategy will either.

We define the best reply correspondence $BR : G \rightarrow G$ player by player as $BR_n : G \rightarrow G_n$ by

$$BR_n(g) = \arg \max_{x \in G_n} u_n(x, g_{-n})$$

and, for every $\varepsilon > 0$, the ε -best reply correspondence $BR^\varepsilon : G \rightarrow G$ player by player as $BR_n^\varepsilon : G \rightarrow G_n$ by

$$BR_n^\varepsilon(g) = \{x \in G_n \mid x(\{b \mid b \notin BR_n(g)\}) \leq \varepsilon\}.$$

And similarly to the earlier analysis we can rather work with \hat{G} so that we have $BR : \hat{G} \rightarrow \hat{G}$ and $BR^\varepsilon : \hat{G} \rightarrow \hat{G}$.

Definition 24. An ε -general form perfect equilibrium is a completely mixed general strategy profile g in \hat{G}^0 such that g is in $BR^\varepsilon(g)$. A general form perfect equilibrium g is a limit of ε -general form perfect equilibria g^ε as $\varepsilon \rightarrow 0$.

The existence of ε -general form perfect equilibria follows from standard arguments since \hat{G} is convex and compact and BR^ε is closed graph nonempty convex valued and $BR^\varepsilon(g) \cap \hat{G}^0 \neq \emptyset$ for all g . And any sequence of ε -general form perfect equilibria has a convergent subsequence since \hat{G} is compact.

REFERENCES

- CARLOS ALÓS-FERRER AND KLAUS RITZBERGER (2016): *The Theory of Extensive Form Games*, Springer Series in Game Theory, Springer-Verlag.
- CARLOS ALÓS-FERRER AND KLAUS RITZBERGER (2017): “Characterizations of Perfect Recall,” *International Journal of Game Theory*, 46(2), 311–326.
- STEVE ALPERN (1988): “Games with Repeated Decisions,” *SIAM Journal of Control and Optimization*, 26, 468–477.
- ROBERT J. AUMANN, SERGIU HART, AND MOTTY PERRY (1997a): “The Absent-Minded Driver,” *Games and Economic Behavior*, 20(1), 102–116.
- ROBERT J. AUMANN, SERGIU HART, AND MOTTY PERRY (1997b): “The Forgetful Passenger,” *Games and Economic Behavior*, 20(1), 117–120.
- PIERPAOLO BATTIGALLI (1997): “Dynamic Consistency and Imperfect Recall,” *Games and Economic Behavior*, 20(1), 31–50.
- LAWRENCE E. BLUME AND WILLIAM R. ZAME (1994): “The Algebraic Geometry of Perfect and Sequential Equilibrium,” *Econometrica*, 62, 783–794.
- ITZHAK GILBOA (1997): “A Comment on the Absent-Minded Driver Paradox,” *Games and Economic Behavior*, 20(1), 25–30.
- ADAM J. GROVE AND JOSEPH Y. HALPERN (1997): “On the Expected Value of Games with Absentmindedness,” *Games and Economic Behavior*, 20(1), 51–65.
- JOSEPH Y. HALPERN (1997): “On Ambiguities in the Interpretation of Game Trees,” *Games and Economic Behavior*, 20(1), 66–96.
- JOSEPH Y. HALPERN AND RAFAEL PASS (2017): “Sequential Equilibrium in Games of Imperfect Recall,” unpublished.
- EBBE HENDON, HANS JØRGEN JACOBSEN, AND BIRGITTE SLOTH (1996): “The One-Shot-Deviation Principle for Sequential Rationality,” *Games and Economic Behavior*, 12(2), 274–282.
- JOHN HILLAS, TINA KAO, AND AARON SCHIFF (2017): “On a Theorem of Blume and Zame,” unpublished.
- JOHN R. ISBELL (1957): “Finitary Games,” in *Contributions to the Theory of Games, Volume 3*, edited by Melvin Drescher, Albert W. Tucker, and Philip Wolfe, 79–96, Princeton University Press, Princeton, NJ.
- JEFFREY JUDE KLINE (2005): “Imperfect Recall and the Relationships Between Solution Concepts in Extensive Games,” *Economic Theory*, 25(3), 703–710.

- ELON KOHLBERG AND JEAN-FRANÇOIS MERTENS (1986): “On the Strategic Stability of Equilibria,” *Econometrica*, 54, 1003–1038.
- DAVID M. KREPS AND ROBERT WILSON (1982): “Sequential Equilibria,” *Econometrica*, 50, 863–894.
- HAROLD W. KUHN (1950): “Extensive Games,” *Proceedings of the National Academy of Sciences*, 36, 570–576.
- HAROLD W. KUHN (1953): “Extensive Games and the Problem of Information,” in *Contributions to the Theory of Games, Volume 2*, edited by Harold W. Kuhn and Albert W. Tucker, 193–216, Princeton University Press, Princeton, NJ.
- BARTON L. LIPMAN (1997): “More Absentmindedness,” *Games and Economic Behavior*, 20(1), 97–101.
- ADRIAN MARPLE AND YOAV SHOHAM (2013): “Equilibria in Finite Games with Imperfect Recall,” unpublished.
- JEAN-FRANÇOIS MERTENS, SYLVAIN SORIN, AND SHMUEL ZAMIR (2015): *Repeated Games*, Econometric Society Monographs, Cambridge University Press, Cambridge.
- ROGER MYERSON (1978): “Refinement of the Nash Equilibrium Concept,” *International Journal of Game Theory*, 7, 73–80.
- ANDRÉS PEREA (2002): “A Note on the One-Deviation Property in Extensive Form Games,” *Games and Economic Behavior*, 40(2), 322 – 338.
- MICHELE PICCIONE AND ARIEL RUBINSTEIN (1997a): “On the Interpretation of Decision Problems with Imperfect Recall,” *Games and Economic Behavior*, 20(1), 3–24.
- MICHELE PICCIONE AND ARIEL RUBINSTEIN (1997b): “The Absent-Minded Driver’s Paradox: Synthesis and Responses,” *Games and Economic Behavior*, 20(1), 121–130.
- CARLOS PIMIENTA AND JIANFEI SHEN (2013): “On the Equivalence Between (Quasi-)Perfect and Sequential Equilibria,” *International Journal of Game Theory*, 43, 1–8.
- KLAUS RITZBERGER (1999): “Recall in Extensive Form Games,” *International Journal of Game Theory*, 28, 69–87.
- REINHARD SELTEN (1965): “Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragetragheit,” *Zeitschrift für die gesamte Staatswissenschaft*, 121, 301–324, 667–689.
- REINHARD SELTEN (1975): “Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games,” *International Journal of Game Theory*, 4, 25–55.
- ERIC VAN DAMME (1984): “A Relation between Perfect Equilibria in Extensive Form Games and Proper Equilibria in Normal Form Games,” *International Journal of Game Theory*, 13, 1–13.
- JOHN VON NEUMANN AND OSKAR MORGENSTERN (1947): *Theory of Games and Economic Behavior*, Princeton University Press, Princeton NJ, 2nd edition.

THE UNIVERSITY OF AUCKLAND, NEW ZEALAND

WASEDA UNIVERSITY, JAPAN