Proceedings, Fifteenth International Conference on
Principles of Knowledge Representation and Reasoning (KR 2016)

Sequential Equilibrium in Games of Imperfect Recall

Joseph Y. Halpern*
Cornell University
halpern@cs.cornell.edu

Rafael Pass†
Cornell University
rafael@cs.cornell.edu

Abstract

Definitions of sequential equilibrium and perfect equilibrium are given in games of imperfect recall. Subtleties regarding the definition are discussed.

1 Introduction

There has been a great deal of interest in AI recently in applying ideas of game theory to model interacting agents who have possibly different preferences as to the outcome of the interaction. While perhaps the best-known solution concept for games is Nash equilibrium, a Nash equilibrium may rely on threats; there might be some concern that these are empty threats, which will never be carried out. The notions of *sequential equilibrium* (Kreps and Wilson 1982) and *perfect equilibrium* (Selten 1975) are perhaps the most common solution concepts used to deal with empty threats. They are both trying to capture the intuition that agents play optimally, not just on the equilibrium path (which is all the is enforced by Nash eqiulibrium), but also off the equilibrium path.

Unfortunately, both perfect equilibrium and sequential equilibrium have been defined only in games of perfect recall, where players remember all the moves that they have made and what they have observed. Perfect recall seems to be an unreasonable assumption in practice. To take just one example, consider even a relatively short card game such as bridge. In practice, in the middle of a game, most people do not remember the complete bidding sequence and the complete play of the cards (although this can be highly relevant information!). Indeed, more generally, we would not expect most people to exhibit perfect recall in games that are even modestly longer than the standard two- or three-move games considered in most game theory papers. Perhaps even more relevant for AI, it is unreasonable to expect software agents to have perfect recall; this requires far too much memory, in general.

Nevertheless, the intuition that underlies sequential and perfect equilibrium, namely, players should play optimally even off the equilibrium path, seems to make sense even in games of imperfect recall. An agent with imperfect recall will still want to play optimally in all situations. And although, in general, calculating what constitutes optimal play may be complicated (indeed, the definition of sequential equilibrium is itself complicated), there are many games where it is not that hard to do. However, the work of Piccione and Rubinstein (1997b) (PR from now on) suggests some subtleties. The following two examples, both due to PR, illustrate the problems.

Example 1.1: Consider the game described in Figure 1, which we call the "match nature" game:

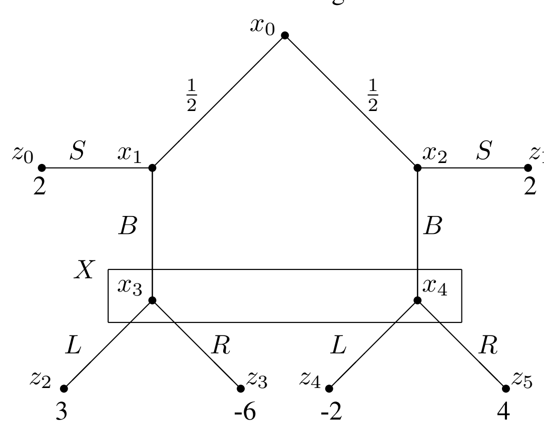

Figure 1: Subtleties with imperfect recall, illustrated by the match nature game.

At x_0, nature goes either left or right, each with probability $1/2$. At x_1, the agent knows that nature moved left; at x_2, the agent knows that nature moved right. The box around the nodes x_3 and x_4 labeled X indicates that $X = \{x_3, x_4\}$ is an *information set*. Thus, at the information set X, the agent

*Supported in part by NSF grants IIS-0534064, IIS-0812045, IIS-0911036, and CCF-1214844, and by AFOSR grants FA9550-08-1-0438, FA9550-09-1-0266, and FA9550-12-1-0040, and ARO grant W911NF-09-1-0281.

†Supported in part by an Alfred P. Sloan Fellowship, a Microsoft New Faculty Fellowship, NSF Awards CNS-1217821 and CCF-1214844, NSF CAREER Award CCF-0746990, AFOSR Award FA9550-08-1-0197, AFOSR YIP Award FA9550-10-1-0093, BSF Grant 2006317, and DARPA and AFRL under contract FA8750-11-2-0211. The views and conclusions contained in this document are those of the authors and should not be interpreted as the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

has forgotten what nature did. Since x_3 and x_4 are in the same information set, the agent must do the same thing at both. It is not hard to show that the strategy that maximizes expected utility chooses action S at node x_1, action B at node x_2, and action R at the information set X consisting of x_3 and x_4. Call this strategy b. Let b' be the strategy of choosing action B at x_1, action S at x_2, and L at X. As PR point out, if node x_1 is reached and the agent is using b, then he will not feel that b is optimal, conditional on being at x_1; he will want to use b'. Indeed, there is *no* single strategy that the agent can use that he will feel is optimal both at x_1 and x_2. ∎

The problem here is that if the agent starts out using strategy b and then switches to b' if he reaches x_1 (but continues to use b if he reaches x_2), he ends up using a "strategy" that does not respect the information structure of the game, since he makes different moves at the two nodes in the information set $X = \{x_3, x_4\}$.[1] As pointed out by Halpern (1997), if the agent knows what strategy he is using at all times, and he is allowed to change strategies, then the information sets are not doing a good job here of describing what the agent knows, since the agent can be using different strategies at two nodes in the same information set. The agent will know different things at x_3 and x_4, despite them being in the same information set.

Example 1.2: The following game, commonly called the *absent-minded driver paradox*, illustrates a different problem. It is described by PR as follows:

> An individual is sitting late at night in a bar planning his midnight trip home. In order to get home he has to take the highway and get off at the second exit. Turning at the first exit leads into a disastrous area (payoff 0). Turning at the second exit yields the highest reward (payoff 4). If he continues beyond the second exit he will reach the end of the highway and find a hotel where he can spend the night (payoff 1). The driver is absentminded and is aware of this fact. When reaching an intersection he cannot tell whether it is the first or the second intersection and he cannot remember how many he has passed.

The situation is described by the game tree in Figure 2.

Clearly the only decision the driver has to make is whether to get off when he reaches an exit. A straightforward computation shows that the driver's optimal strategy *ex ante* is to exit with probability $1/3$; this gives him a payoff of $4/3$. On the other hand, suppose that the driver starts out using the optimal strategy, and when he reaches the information set, he ascribes probability α to being at e_1. He then considers whether he should switch to a new strategy, where he exits with probability p. Another straightforward calculation

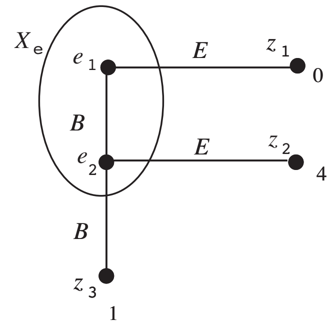

Figure 2: The absentminded driver game.

shows that his expected payoff is then

$$\begin{aligned} &\alpha((1-p)^2 + 4p(1-p)) + (1-\alpha)((1-p) + 4p) \\ &= 1 + (3-\alpha)p - 3\alpha p^2. \end{aligned} \tag{1}$$

Equation 1 is maximized when $p = \min(1, (3-\alpha)/6\alpha)$, with equality holding only if $\alpha = 1$. Thus, unless the driver ascribes probability 1 to being at e_1, he should want to change strategies when he reaches the information set. This means that as long as $\alpha < 1$, we cannot hope to find a sequential equilibrium in this game. The driver will want to change strategies as soon as he reaches the information set. ∎

According to the technique used by Selten (1975) to ascribe beliefs, also adopted by PR, if the driver is using the optimal strategy, e_1 should have probability $3/5$ and e_2 should have probability $2/5$. The argument is that, according to the optimal strategy, e_1 is reached with probability 1 and e_2 is reached with probability $2/3$. Thus, 1 and $2/3$ should give the relative probability of being at e_1 and e_2. Normalizing these numbers gives us $3/5$ and $2/5$, and leads to non-existence of sequential equilibrium. (This point is also made by Kline (2005).)

As shown by PR and Aumann, Hart, and Perry (AHP) (1997), this way of ascribing beliefs guarantees that the driver will not want to use any *single-action deviation* from the optimal strategy. That is, there is no "strategy" b' that is identical to the optimal strategy except at one node and has a higher payoff than the optimal strategy. PR call this the *modified multi-self approach*, whereas AHP call it *action-optimality*. AHP suggest that this approach solves the paradox. On the other hand, Piccione and Rubinstein (1997a) argue that it is hard to justify the assumption that an agent cannot change her future actions. (See also (Gilboa 1997; Lipman 1997) for further discussion of this issue.)

Our goal in this paper is to define sequential equilibrium equilibrium for games of imperfect recall. As these examples show, such definitions require a clear interpretation of the meaning of information sets and the restrictions they impose on the knowledge and strategies of players. Moreover, as we shall show, there are different intuitions behind the notion of sequential equilibrium. While they all lead to the same definition in games of perfect recall, this is no longer

[1] As usual, we take a pure strategy b to be a function that associates with each node in the game tree a move, such that if x and x' are two nodes in the same information set, then $b(x) = b(x')$. We occasionally abuse notation and write "strategy" even for a function b' that does not necessarily satisfy the latter requirement; that is, we may have $b'(x) \neq b'(x')$ even if x and x' are in the same information set.

the case in games of imperfect recall. Our definition can be viewed as trying to capture a notion of *ex ante* sequential equilibrium. The picture here is that players choose their strategies before the game starts and are committed to it, but they choose it in such a way that it remains optimal even off the equilibrium path. This, unfortunately, does *not* correspond to the more standard intuitions behind sequential equilibrium, where agents are reconsidering at each information set whether they are doing the "right" thing, and can change their strategies if they should choose to. While we believe that defining such a notion of *interim* sequential rationality would be of great interest (and discuss potential definitions of such a notion in the full paper), it raises a number of new subtleties in games of incomplete information, since the obvious definition is in general incompatible with the exogenously-given information structure. (This is essentially the point made by the match nature game in Figure 1.) We believe that having a definition of sequential rationality that agrees with the standard definition of games of perfect recall and is conceptually clear and well motivated in games of imperfect recall will help us understand better the interplay between rationality and imperfect recall. Moreover, the *ex ante* notion is particularly well motivated in a setting where players are choosing an algorithm, and are charged for the complexity of the algorithm, in the spirit of the work of Rubinstein (1986); we explore some of these issues in definitions of sequential equilibrium in this setting, based on the ideas of this paper, in (Halpern and Pass 2013).

2 Preliminaries

In this section, we discuss a number of issues that will be relevant to our definition of sequential equilibrium: imperfect recall and absentmindedness, what players know, behavioral vs. mixed strategies, and belief ascription.

2.1 Imperfect Recall and Absentmindedness

As usual, we describe an *extensive-form game*, where decisions are made over time, using a game tree and information sets (as in Figures 1 and 2). A game is said to exhibit *perfect recall* if, for all players i and all nodes x_1 and x_2 in an information set X for player i, if h_j is the history leading to x_j for $j = 1, 2$, player i has played the same actions in h_1 and h_2 and gone through the same sequence of information sets. (See, for example, (Osborne and Rubinstein 1994) for more formal definitions.) If a game does not exhibit perfect recall, it is said to be a game of *imperfect recall*. A special case of imperfect recall is *absentmindedness*; absentmindedness occurs when there are two nodes on one history that are in the same information set. The absent-minded driver game exhibits absentmindedness; the match nature game does not.

2.2 Knowledge of Strategies

The standard (often implicit) assumption in most game theory papers is that players know their strategies. This is a nontrivial assumption. Consider the match nature game. If the agent cannot recall his strategy, then certainly any discussion of reconsideration at x_2 becomes meaningless; there is no reason for the agent to think that he will realize at x_4 that he should play R. But if the agent cannot recall even his initial choice of strategy (and thus cannot commit to a strategy) then strategy b (playing B at x_1, S at x_2, and R at X) may not turn out to be optimal. When the agent reaches S, he may forget that he was supposed to play R. While it could be argued that, as long as the agent remembers the structure of the game, he can recompute the optimal strategy. However, this argument no longer holds if we change the payoffs at z_4 and z_5 to -6 and 3, respectively, so that the left and right sides of the game tree are completely symmetric. Then it is hard to see how an agent who does not recall what strategy he is playing will know whether to play L or R at X. A prudent agent might well decide to play S at both x_1 and x_2! Because we are considering an *ex ante* notion of sequential equilibrium, we assume that an agent can commit initially to playing a strategy (and will know this strategy at later nodes in the game tree). But we stress that we view this assumption as problematic if we allow reconsideration of strategies at later information sets.

2.3 Mixed Strategies vs. Behavioral Strategies

There are two types of strategies that involve randomization that have been considered in extensive-form games. A *mixed strategy* in an extensive-form game is a probability measure on pure strategies. Thus, we can think of a mixed strategy as corresponding to a situation where a player tosses a coin and chooses a pure strategy at the beginning of the game depending on the outcome of the coin toss, and then plays that pure strategy throughout the game. By way of contrast, with a *behavioral* strategy, a player randomizes at each information set, randomly choosing an action to play at that information set. Formally, a behavioral strategy is a function from information sets to distributions over acts. (We can identify a pure strategy with the special case of a behavioral strategy that places probability 1 on some action at every information set.) Thus, we can view a behavioral strategy for player i as a collection of probability measures indexed by the information sets for player i; there is one probability measure on the actions that can be performed at information set X for each information set X for player i.

It is well known that in games of perfect recall, mixed strategies and behavioral strategies are *outcome-equivalent*. That is, given a mixed strategy b for player i, there exists a behavioral strategy b' such that, no matter what strategy profile (mixed or behavioral) b_{-i} the remaining players use, (b, b_{-i}) and (b', b_{-i}) induce the same distribution on the leaves (i.e., terminal histories) of the game tree; and conversely, for every mixed strategy b, there exists a behavioral strategy b' such that for all strategy profiles b_{-i} for the remaining player, (b, b_{-i}) and (b', b_{-i}) are outcome-equivalent. (See (Osborne and Rubinstein 1994) for more details.)

It is also well known that this equivalence breaks down when we move to games of imperfect recall. In games without absentmindedness, for every behavioral strategy, there is an outcome-equivalent mixed strategy; but, in general, the converse does not hold (Isbell 1957). Once we allow absentmindedness, as pointed out by PR, there may be behavioral strategies that are not outcome-equivalent to any mixed strategy. This is easy to see in the absentminded driver game. The two pure strategies reach z_1 and z_3, respectively. Thus, no mixed strategy can reach z_2, while any behavioral strategy that places positive probability on both B and E has some positive probability of reaching z_2. The same argument also shows that there exist mixed strategies that are not outcome-equivalent to any behavioral strategy.

Thus, to deal with games of imperfect recall, in general, we need to allow what Kaneko and Kline (1995) call *behavioral strategy mixtures*.[2] A behavioral strategy mixture b_i for player i is a probability distribution on behavioral strategies for player i that assigns positive probability to only finitely many behavioral strategies for player i. As Kaneko and Kline note, a behavioral strategy mixture involves two kinds of randomization: before the game and in the course of the game. A behavioral strategy is the special case of a behavioral strategy mixture where the randomization happens only during the course of the game; a mixed strategy is the special case where the randomization happens only at the beginning. For the remainder of the paper, when we say "strategy", we mean "behavioral strategy mixture", unless we explicitly say otherwise.

It is worth noting that players do not have to mix over too many behavioral strategies when employing a behavioral strategy mixture. Specifically, we show that a behavioral strategy mixture for player i in a game Γ is outcome-equivalent to a mixture of at most $d_i + 1$ behavioral strategies, where d_i is a constant that depends only of the size of the game Γ. A behavioral strategy mixture for a player i in a game Γ can be described by specifying for each information set I for i a probability distribution over nodes $x' \notin I$ that can be reached from some node $x \in I$. If we set d_i to be the total number of *final nodes* of any information set I for i (where x is a final node of I if there are no nodes x' in I that come after x in the game tree), then by Carathéodory's Theorem (Rockafellar 1970), which says that any point in the convex hull of a set P in R^d is the convex combination of at most $d + 1$ points in P, it follows that a behavioral strategy mixture for player i in a game Γ is outcome-equivalent to a mixture of at most $d_i + 1$ behavioral strategies.[3]

[2] They actually call them *behavior strategy mixtures*; we write "behavioral" for consistency with our terminology elsewhere.

[3] In a game of perfect recall, any behavioral strategy is outcome-equivalent to a convex combination of pure strategies, so there is a fixed finite set P such that every behavioral strategy is a convex combination of the strategies in P. But in a game of imperfect recall, although there is a fixed d such that every behavioral strategy mixture is convex combination of $d + 1$ behavioral strategies, there is no finite set P such that every behavioral strategy mixture is a convex combination of the behavioral strategies in P. (Proof sketch: Consider the absentminded driver game. Behavioral strategies lead to a distribution over leaves of the form $(x, x(1-x), (1-x)^2)$, for $x \in [0, 1]$. Thus, if we view the distribution as a vector of the form (a, b, c), we must have $b \leq a$, with $b = a$ iff $a = b = 0$, and we can make the ratio of b/a as close to 1 as we like, by making a sufficiently small. Any finite collection P of behavioral strategies has a maximum value for b/a. Thus, for any finite set P of behavioral strategies in the absentminded driver game, there must be a behavioral strategy that is not in the convex closure of P.)

A consequence of this fact is that we can identify a mixed behavioral strategy for player i with an element of $([0, 1] \times R^{d_i})^{(d_i+1)}$—each mixed behavioral strategy can be viewed as a tuple of the form $((a_1, b_1), ..., (a_{d_i+1}, b_{d_i+1}))$, where $a_1, \ldots, a_{d_i+1} \in [0, 1]$, $\sum a_i = 1$, and b_j is a mixed behavioral strategy for player i, and thus in R^{d_i}. Since it is well known that the convex closure of a compact set in a finite-dimensional space is closed (Rockafellar 1970), it follows that the set of behavioral strategy mixtures of a finite game Γ is closed, and thus also compact. Summarizing this discussion, we have the following proposition.

Proposition 2.1: *If Γ is a finite game, then the set of behavioral strategy mixtures in Γ is compact. Furthermore, there exists some constant d (which depends on the number of actions in Γ) such that for every behavioral strategy mixture c, there exists an outcome-equivalent behavioral strategy mixture c' that mixes only over d behavioral strategies.*

Solution concepts typically depend only on outcomes, and so are insensitive to the replacement of strategies by outcome-equivalent strategies. For example, if a strategy profile b is a Nash (resp., sequential) equilibrium, and b_i is outcome-equivalent to b'_i, then b' is also a Nash (resp., sequential) equilibrium. Nash showed that every finite game has a Nash equilibrium in mixed strategies. By the outcome-equivalence mentioned above, in a game of perfect recall, there is also a Nash equilibrium that is a behavioral strategy profile. This is no longer the case in games of imperfect recall (Wichardt 2008).

Sequential equilibrium is usually defined in terms of behavioral strategies. This is because it is typically presented as an interim notion. That is, players are viewed as making decision regarding whether they should change strategies at each information set. Thus, it makes sense to view them as using behavioral strategies rather than mixed strategies. Although we view our notion of sequential equilibrium as an *ex ante* notion, we allow agents to use behavioral strategy mixtures. The interpretation is that the agent randomizes at the beginning to choose a behavioral strategy (one that is compatible with the information structure of the game). The agent then commits to this behavioral strategy and follows it throughout the game. The agent has the capability to randomize at each information set, but he is committed to doing the randomization in accordance with his *ex ante* choice.

2.4 Expected Utility of Strategies

Every behavioral strategy mixture profile b induces a probability measure π_b on leaves (terminal histories). We identify a node x in a game with the event consisting of the leaves that can be reached from x. In the language of Grove and Halpern (1997), we are identifying x with the event of *reaching x*. Given this identification, we take $\pi_b(x)$ to be the probability of reaching a leaf that comes after x when using strategy b.

For the purposes of this discussion, fix a game Γ, and let Z denote the leaves (i.e., terminal histories) of Γ. As usual, we can take $EU_i(b)$ to be $\sum_{z \in Z} \pi_b(z) u_i(z)$. If Y is a subset of leaves such that $\pi_b(Y) > 0$, then computing the expected utility of b for player i conditional on Y is equally straightforward. It is simply $EU_i(b \mid Y) = \sum_{z \in Y} \pi_b(z \mid Y) u_i(z)$.

3 Beliefs in Games of Imperfect Recall

Fix a game Γ. Following Kreps and Wilson (1982), a *belief system* μ for Γ is a function that associates with each information set X in Γ a probability μ_X on the histories in X. PR quite explicitly interpret $\mu_X(x)$ as the probability of *being at* the node x, conditioned on reaching X. Just as Kreps and Wilson, they thus require that $\sum_{x \in X} \mu_X(x) = 1$.

Since we aim to define an *ex ante* notion of sequential rationality, we instead interpret $\mu_X(x)$ as the probability of *reaching* x, conditioned on reaching X. We no longer require that $\sum_{x \in X} \mu_X(x) = 1$. While this property holds in games of perfect recall, in games of imperfect recall, if X contains two nodes that are both on a history that is played with positive probability, the sum of the probabilities will be greater than 1. For instance, in the absent minded driver's game, the *ex ante* optimal strategy reaches e_1 with probability 1 and reaches e_2 with probability $2/3$.

Given an information set X, let the *upper frontier of* X (Halpern 1997), denoted $\hat{X}$, to consist of all those nodes $x \in X$ such that there is no node $x' \in X$ that strictly precedes x on some path from the root. Note that for games where there is at most one node per history in an information set, we have $\hat{X} = X$. Rather than requiring that $\sum_{x \in X} \mu_X(x) = 1$, we require that $\sum_{x \in \hat{X}} \mu_X(x) = 1$, that is, that the probability of reaching the upper-frontier of X, conditional on reaching X, be 1. Since $\hat{X} = X$ in games of perfect recall, this requirement generalizes that of Kreps and Wilson. Moreover, it holds if we define μ_X in the obvious way.

Claim 3.1: If X is an information set that is reached by strategy profile b with positive probability and $\mu_X(x) = \pi_b(x \mid X)$, then $\sum_{x \in \hat{X}} \mu_X(x) = 1$.

Proof: By definition, $\sum_{x \in \hat{X}} \mu_X(x) = \sum_{x \in \hat{X}} \pi_b(x \mid X) = \pi(\hat{X} \mid X) = 1$. ∎

Given a belief system μ and a strategy profile b, define a probability distribution μ_X^b over terminal histories in the obvious way: for each terminal history z, if there is no prefix of z in X, then $\mu_X^b(z) = 0$; otherwise, let x_z be the shortest history in X that is a prefix of z, and define $\mu_X^b(z)$ to be the product of $\mu_X(x_z)$ and the probability that b leads to the terminal history z when started in x_z.

Our definition of the probability distribution μ_X^b induced over terminal histories is essentially equivalent to that used by Kreps and Wilson (1982) for games of perfect recall. The only difference is that in games of perfect recall, a terminal history has at most one prefix in X. This no longer is the case in games of imperfect recall, so we must specify which prefix to select. For definiteness, we have taken the shortest prefix; however, it is easy to see that if $\mu_X(x)$ is defined as the probability of b reaching x conditioned on reaching X, then any choice of x_z leads to the same distribution over terminal histories (as long as we choose consistently, that is, we take $x_z = x_{z'}$ if z and z' have a common ancestor in X).

Note that if a terminal history z has a prefix in X, then the shortest prefix of z in X is in $\hat{X}$. Moreover, defining μ_X^b in terms of the shortest history guarantees that it is a well-defined probability distribution, as long as $\sum_{x \in \hat{X}} \mu_X(x) = 1$, even if μ_X is not defined by some strategy profile b'.

Claim 3.2: If Z is the set of terminal histories and $\sum_{x \in \hat{X}} \mu_X(x) = 1$, then for any strategy profile b, we have $\sum_{z \in Z} \mu_X^b(z) = 1$.

Proof: By definition,

$$\begin{array}{rcl} \sum_{z \in Z} \mu_X^b(z) & = & \sum_{z \in Z} \mu_X(x_z) \pi_b(z \mid x_z) \\ & = & \sum_{z \in Z} \sum_{x \in \hat{X}} \mu_X(x) \pi_b(z \mid x) \\ & = & \sum_{x \in \hat{X}} \mu_X(x) \sum_{z \in Z} \pi_b(z \mid x) \\ & = & \sum_{x \in \hat{X}} \mu_X(x) \\ & = & 1. \end{array}$$

∎

Following Kreps and Wilson (1982), let $EU_i(b \mid X, \mu)$ denote the expected utility for player i, where the expectation is taken with respect to μ_X^b.

The following proposition justifies our method for ascribing beliefs. Say that a belief system μ is *compatible* with a strategy b if, for all information sets X such that $\pi_b(X) > 0$, we have $\mu_X(x) = \pi_b(x \mid X)$.

Proposition 3.3: *If information set X is reached by strategy profile b with positive probability, and μ is compatible with b, then*

$$EU_i(b \mid X) = EU_i(b \mid X, \mu).$$

Proof: Let Z be the set of terminal histories, and let Z_X consist of those nodes in Z that have a prefix in X; similarly, let Z_x consist of those nodes in Z whose prefix is x. Using the fact that $\hat{X} = \{x_z : z \in Z_X\}$, we get that

$$\begin{array}{rl} & EU_i(b \mid X, \mu) \\ = & \sum_{z \in Z} \mu_X^b(z) u_i(z) \\ = & \sum_{z \in Z_X} \mu_X(x_z) \pi_b(z \mid x_z) u_i(z) \\ = & \sum_{z \in Z_X} \pi_b(x_z \mid X) \pi_b(z \mid x_z) u_i(z) \\ = & \sum_{z \in Z_X} \pi_b(z \mid X) u_i(z) \\ = & EU_i(b \mid X). \end{array}$$

∎

4 Perfect and Sequential Equilibrium

4.1 Perturbed Games

Weak compatibility tells us how to define beliefs for information sets that are on the equilibrium path. But it does not tell us how to define the beliefs for information sets that are off the equilibrium path. We need to know how to do this in order to define both sequential and perfect equilibrium. To deal with this, we follow Selten's approach of considering *perturbed games*. Given an extensive-form game Γ and a function η associating with every information set X and action c that can be performed at X, a probability $\eta_c > 0$ such that, for each information set X for player i, if $A(X)$ is the set of actions that i can perform at X, then $\sum_{c \in A(X)} \eta_c < 1$.

We call η a *perturbation* of Γ. We think of η_c as the probability of a "tremble"; since we view trembles as unlikely, we are most interested in the case that η_c is small.

A *perturbed game* is a pair (Γ, η) consisting of a game Γ and a perturbation η. A behavioral strategy b for player i in (Γ, η) is *acceptable* if, for each information set X and each action $c \in A(X)$, $b(X)$ assigns probability at least η_c to c. A behavioral strategy mixture b is acceptable in (Γ, η) if for each information set X and each action $c \in A(X)$, the expected probability of playing c according to b is at least η_c.

We can define best responses and Nash equilibrium in the usual way in perturbed games (Γ, η); we simply restrict the definitions to the acceptable strategies for (Γ, η). Note that if b is an acceptable strategy profile in a perturbed game, then $\pi_b(X) > 0$ for all information sets X.

4.2 Best Responses at Information Sets

There are a number of ways to capture the intuition that a strategy b_i for player i is a best response to a strategy profile b_{-i} for the remaining players at an information set X. To make these precise, we need some notation. Given a behavioral strategy b, let $b_i[X/c]$ denote the behavioral strategy that is identical to b_i except that, at information set X, action c is played.

Switching to another action at an information set is, of course, not the same as switching to a different strategy at an information set. If b' is a strategy for player i, we would like to define the strategy $[b_i, X, b']$ to be the strategy where i plays b up to X, and then switches to b' at X. Intuitively, this means that i plays b' at all information sets that come after X. The problem is that the meaning of "after X" is not so clear in games with imperfect recall. For example, in the match nature game, is the information set X after the information set $\{x_1\}$? While x_3 comes after x_1, x_4 does not. The obvious interpretation of switching from b to b' at x_1 would have the agent playing b' at x_3 but still using b at x_4. As we have observed, the resulting "strategy" is not a strategy in the game, since the agent does not play the same way at x_3 and x_4.

This problem does not arise in games of perfect recall. Define a strict partial order $\prec$ on nodes in a game tree by taking $x \prec x'$ if x precedes x' in the game tree. There are two ways to extend this partial order on nodes to a partial order on information sets. Given information sets X and X' for a player i, define $X \prec X'$ iff, for all $x' \in X'$, there exists some $x \in X$ such that $x \prec x'$. It is easy to see that $\prec$ is indeed a partial order. Now define $X \prec' X'$ iff, for *some* $x' \in X'$ and $x \in X$, $x \prec x'$. It is easy to see that $\prec$ agrees with $\prec'$ in games of perfect recall. However, they do *not* in general agree in games of imperfect recall. For example, in the match nature game, $\{x_2\} \prec' X$, but it is not the case that $\{x_2\} \prec X$. Moreover, although $\prec'$ is a partial order in the match nature game, in general, in games of imperfect recall, $\prec'$ is not a partial order. In particular, in the game in Figure 3, we have $X_1 \prec' X_2$ and $X_2 \prec' X_1$.

Define $x \preceq x'$ iff $x = x'$ or $x \prec x'$; similarly, $X \preceq X'$ iff $X = X'$ or $X \prec X'$. We can now define $[b, X, b']$ formally, where b is a behavioral strategy mixture and b' is a behavioral strategy.

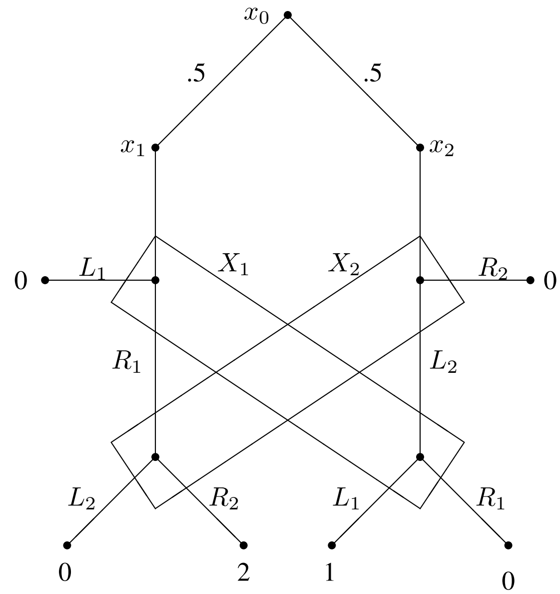

Figure 3: A game where $\prec'$ is not a partial order.

We start by defining $[b, X, b']$ when both b and b' are behavioral strategies. In that case, $[b, X, b']$ is the strategy according to which i plays b' at every information set X' such that $X \preceq X'$, and otherwise plays b. If c is a behavioral strategy mixture and b' is a behavioral strategy then $[c, X, b']$ is the behavioral strategy mixture that puts probability $c(b)$ on the behavioral strategy $[b, X, b']$. We do not define $[c, X, c']$ in case c' is a behavioral strategy mixture (as opposed to just a behavioral strategy); randomization over behavioral strategies is allowed only at the beginning of the game.

The strategy $[b, X, b']$ is well defined even in games of imperfect recall, but it is perhaps worth noting that the strategy $[b, \{x_1\}, b']$ in the match nature game is the strategy where the player goes down at x_1, but still plays R at information set X, since we do not have $X \preceq \{x_1\}$. Thus, $[b, \{x_1\}, b']$ as we have defined it is not better for the player than b. Similarly, in the game in Figure 3, if b is the strategy of playing R_1 at X_1 and R_2 at X_2, while b' is the strategy of playing L_1 at X_1 and L_2 at X_2, then $[b, \{x_2\}, b']$ is b.

If we are thinking in terms of players switching strategies, then strategies of the form $[b, X, b']$ allow as many switches as possible. To make this more precise, if b and b' are behavioral strategies, let (b, X, b') denote the "strategy" of using b until X is reached and then switching to b'. More precisely, $(b, X, b')(x) = b'(x)$ if $x' \prec x$ for some node $x' \in X$; otherwise, $(b, X, b')(x) = b(x)$. Intuitively, (b, X, b') switches from b to b' as soon as a node in X is encountered. As observed above, (b, X, b') is not always a strategy. But whenever it is, $(b, X, b') = [b, X, b']$.

Proposition 4.1: *If c is a behavioral strategy mixture and b' is behavioral strategy, then (c, X, b') is a strategy in game Γ iff $(c, X, b') = [c, X, b']$.*

Proof: We first prove the proposition in case c is a pure behavioral strategy b. Suppose that $(b, X, b') \neq [b, X, b']$.

Then there must exist some information set X' such that (b, X, b') and $[b, X, b']$ differ at X'. If $X \preceq X'$, then at every node in $x \in X'$, the player plays $b'(X')$ at x according to both (b, X, b') and $[b, X, b']$. Thus, it must be the case that $X \not\preceq X'$. This means the player plays $b(X')$ at every node in X' according to $[b, X, b']$. Since (b, X, b') and $[b, X, b']$ disagree at X', it must be the case that the player plays $b'(X')$ at some node $x \in X'$ according to (b, X, b'). But since $X \not\preceq X'$, there exists some node $x' \in X'$ that does not have a prefix in X. This means that (b, X, b') must play $b(X')$ at x'. Thus, (b, X, b') is not a strategy.

Now suppose that c is a nontrivial mixture over behavioral strategies, and that $(c, X, b') \neq [c, X, b']$. Thus, there exists some b in the support of c such that $(b, X, b') \neq [b, X, b']$; by the argument above, (b, X, b') is not a strategy, so (c, X, b') cannot be one either.

Conversely, if $(c, X, b') = [c, X', b]$, then clearly (c, X, b') is a strategy. ∎

4.3 Defining Perfect and Sequential Equilibrium

We now define (our versions of) perfect and sequential equilibrium for games of imperfect recall.

Perfect equlibrium: We start with perfect equilibrium. Here we use literally the same definition as Selten (1975), except that we use behavioral strategy mixtures rather than behavioral strategies. The strategy profile b^* is a *perfect equilibrium* in Γ if there exists a sequence $(\Gamma, \eta_1), (\Gamma, \eta_2), \ldots)$ and a sequence of behavioral strategy mixture profiles $b^1, b^2, \ldots$ such that (1) $\eta_k \rightarrow \vec{0}$; (2) b^k is a Nash equilibrium of (Γ, η_k); and (3) $b^k \rightarrow b^*$.[4] Selten (1975) shows that a perfect equilibrium always exists in games with perfect recall. Essentially the same proof shows that it exists even in games with imperfect recall.

Theorem 4.2 : *A perfect equilibrium exists in all finite games.*

Proof: Consider any sequence $(\Gamma, \eta_1), (\Gamma, \eta_2), \ldots)$ of perturbed games such that $\eta_n \rightarrow \vec{0}$. By standard fixed-point arguments, each perturbed game (Γ, η_k) has a Nash equilibrium b^k in behavioral strategy mixtures. (Here we are using the fact that, by Proposition 2.1, the set of behavioral strategy mixtures is compact.) By a standard compactness argument, the sequence $b^1, b^2, \ldots$ has a convergent subsequence. Suppose that this subsequence converges to b^*. Clearly b^* is a perfect equilibrium. (Although the behavioral strategy mixtures in the profile b^* may not have finite support, by Proposition 2.1, we can assume without loss of generality that the mixtures in fact have finite support.) ∎

As we have observed, as a technical matter, using mixed strategies rather than behavioral strategies makes no difference in games of perfect recall. However, it has a major impact in games of imperfect recall. Since it is easy to see that every perfect equilibrium is a Nash equilibrium, it follows from Wichardt's (2008) example that perfect equilibrium does not exist in games of imperfect recall if we restrict to behavioral strategies.

Recall that we view the players as choosing a behavioral strategy mixture at the beginning of the game. They then do the randomization, and choose a behavioral strategy appropriately. At this point, they commit to the behavioral strategy chosen, remember it throughout the game, and cannot change it. However, they make this initial choice in a way that it is not only unconditionally optimal (which is all that is required of Nash equilibrium), but continues to be optimal conditional on reaching each information set.

It is easy to see that a perfect equilibrium b^* of Γ is also a Nash equilibrium of Γ. Thus, each strategy b_i^* is a best response to b_{-i}^* *ex ante*. However, we also want b_i^* to be a best response to b_{-i}^* at each information set. This intuition is made precise using intuitions from the definition of sequential equilibrium, which we now review.

A behavioral strategy b_i is *completely mixed* if, for each information set X and action $c \in A(X)$, b_i assigns positive probability to playing c. A behavioral strategy mixture is completely mixed if every behavioral strategy in its support is completely mixed. A belief system μ is *consistent* with a strategy b if there exists a sequence of completely mixed strategy profiles $b^1, b^2, \ldots$ converging to b such that $\mu_X(x)$ is $\lim_{n\rightarrow\infty} \pi_{b^n}(x \mid X)$. Note that if μ is consistent with b, then it is compatible with b.

The following result makes precise the sense in which a perfect equilibrium is a best response at each information set.

Proposition 4.3: *If b is a perfect equilibrium in game Γ, then there exists a belief system μ consistent with b such that, for all players i, all information sets X for player i, and all behavioral strategies b' for player i, we have*

$$EU_i(b \mid X, \mu) \geq EU_i(([b_i, X, b'], b_{-i}) \mid X, \mu).$$

Proof: Since b is a perfect equilibrium, there exists a sequence of strategy profiles $b^1, b^2, \ldots$ converging to b and a sequence of perturbed games $(\Gamma, \eta_1), (\Gamma, \eta_2), \ldots$ such that $\eta_k \rightarrow \vec{0}$ and b^k is a Nash equilibrium of (Γ, η_k). All the strategies $b^1, b^2, \ldots$ are completely mixed (since they are strategies in perturbed games). We can assume without loss of generality that, for each information set X and $x \in X$, the limit $\lim_{n\rightarrow\infty} \pi_{b^n}(x \mid X)$ exists. (By standard arguments, since Γ is a finite game, by Proposition 2.1, we can find a subsequence of $b^1, b^2, \ldots$ for which the limits all exist, and we can replace the original sequence by the subsequence.) Let μ be the belief assessment determined by this sequence of strategies.

We claim that the result holds with respect to μ. For suppose not. Then there exists a player i, information set X, behavioral strategy b' for player i, and $\epsilon > 0$ such that $EU_i(b \mid X, \mu) + \epsilon < EU_i(([b_i, X, b'], b_{-i}) \mid X, \mu))$. It follows from Proposition 3.3 that $EU_i(b^k \mid X) \rightarrow EU_i(b \mid X, \mu)$ and $EU_i(([b_i^k, X, b'], b_{-i}^k) \mid X) \rightarrow EU_i(([b_i, X, b'], b_{-i}) \mid X, \mu)$. Since $b^k \rightarrow b$ and $\eta_k \rightarrow \vec{0}$, there exists some strategy b'' and $k > 0$ such that b'' is acceptable for

[4] We can work with mixed strategies instead of behavioral strategy mixtures if we do not allow absentmindedness; with absentmindedness, we need behavioral strategy mixtures to get our results.

$(\Gamma, \eta_{k'})$ for all $k' > k$, and $EU_i(b^{k'} \mid X) + \epsilon/2 < EU_i(([b_i^{k'}, X, b''']), b_{-i}^k) \mid X)$. But this contradicts the assumption that $b^{k'}$ in a Nash equilibrium of $(\Gamma, \eta_{k'})$.[5] ∎

We are implicitly identifying "b is a best response for i at information set X" with "$EU_i(b \mid X, \mu) \geq EU_i(([b_i, X, b'], b_{-i}) \mid X, \mu))$" for all behavioral strategies b'. How reasonable is this? In games of perfect recall, if an action at a node x' can affect i's payoff conditional on reaching X, then x' must be in some information set X' after X. This is not in general the case in games of imperfect information. For example, in the match nature game, the player's action at x_4 can clearly affect his payoff conditional on reaching x_2, but the information set X that contains x_4 does not come after $\{x_2\}$, so we do not allow changes at x_4 in considering best responses at x_2. While making a change at x_4 makes things better at x_2, it would make things worse at x_1, a node that is not after x_2. Given our *ex ante* viewpoint, this is clearly a relevant consideration. What we are really requiring is that b is a best response for i at X *among strategies that do not affect i's utility at nodes that do not come after X*. This last phrase does not have to be added in games of perfect recall, but it makes a difference in games of imperfect recall. We return to this point in Section 4.4.

Sequential′ equilibrium: We can now define a notion of sequential equilibrium just as Kreps and Wilson (1982) did. However, this notion turns out not out not to imply Nash equilibrium, so we call it sequential′ equilibrium. We later define a strengthening that we call sequential equilibrium that is better behaved (and arguably better motivated).

Definition 4.4: A *belief assessment* is pair (b, μ) consisting of a strategy b and a belief system μ. A belief assessment (b, μ) is a *sequential′ equilibrium* in a game Γ if μ is consistent with b and, for all players i, all information sets X for player i, and all behavioral strategies b' for player i at X, we have $EU_i(b \mid X, \mu) \geq EU_i(([b_i, X, b'], b_{-i}) \mid X, \mu))$. ∎

It is immediate from Proposition 4.3 that every perfect equilibrium is a sequential′ equilibrium. Thus, a sequential′ equilibrium exists for every game.

Theorem 4.5: *A sequential′ equilibrium exists in all finite games.*

Note that in both Examples 1.1 and 1.2, the *ex ante* optimal strategy is a sequential′ equilibrium according to our definition. In Example 1.1, it is because the switch to what appears to be a better strategy at x_1 is disallowed. In Example 1.2, the unique belief μ consistent with the *ex ante* optimal strategy assigns probability 1 to reaching e_1 and probability $2/3$ to reaching e_2. However, since e_2 is not on the upper frontier of X_e, for all strategies b, $EU(b \mid X_e) = EU(b \mid e_1) = EU(b)$, and thus the *ex ante* optimal strategy is still optimal at X_e.

[5]Note that we here rely on the fact that b' is a behavioral strategy and not a behavioral strategy mixture; if it were a behavioral strategy mixture, then we could no longer guarantee that the "strategy" obtained by switching from b to b' at X is a behavioral strategy mixture (since mixing now happens twice during the game).

Although our definition of sequential′ equilibrium agrees with the traditional definition of sequential equilibrium (Kreps and Wilson 1982) in games of perfect recall, there are a number of properties of sequential equilibrium that no longer hold in games of imperfect recall. First, it is no longer the case that every sequential′ equilibrium is a Nash equilibrium. For example, in the match nature game, it is easy to see that the strategy b' is a sequential′ equilibrium but is not a Nash equilibrium. It is easy to show that every sequential′ equilibrium is a Nash equilibrium in games where each agent has an initial information set that precedes all other information sets (in the $\prec$ order defined above). At such an information set, the agent can essentially do *ex ante* planning. There is no such initial information set in the match nature game, precluding such planning. If we want to allow such planning in a game of imperfect recall, we must model it with an initial information set for each agent.

Summarizing this discussion, we have the following result.

Theorem 4.6: *In all finite games,*

(a) every perfect equilibrium is a Nash equilibrium;

(b) in games where all players have an initial information set, every sequential′ equilibrium is a Nash equilibrium.

However, there exist games where a sequential′ equilibrium is not a Nash equilibrium.

It is also well known that in games of perfect recall, we can replace the standard definition of sequential equilibrium by one where we consider only single-action deviations (Hendon, Jacobsen, and Sloth 1996); this is known as the *one-step deviation principle*. This no longer holds in games of imperfect recall either. Again, consider the modification of the match nature game with an initial node x_{-1}. As we observed above, in this case, starting at x_{-1} by playing down and then playing b is the only sequential′ equilibrium. However, replacing b by b' gives a strategy that satisfies the one-step deviation property.

4.4 Sequential Equilibrium

As we argued above (see the discussion after Proposition 4.3), the sense in which our *ex ante* notions of perfect equilibrium and sequential′ equilibrium capture optimality is that there is (from the *ex ante* point of view) no information set X at which an agent can do better *without affecting his utility at nodes that do not come after X*. This suggests that we might consider a stronger optimality requirement. Instead of looking at just one information set, we can consider a set $\mathcal{X} = \{X_1, \ldots, X_n\}$ of information sets for player i, and require that i not be able to do better at all information sets in $\mathcal{X}$ without affecting his utility at nodes that do not come after $\mathcal{X}$. In games of perfect recall, looking at a set $\mathcal{X}$ of information sets rather than just a single information set does not affect the notion of sequential′ equilibrium. But it does in the case of imperfect recall. Consider the match nature game again. As we observed, the strategy b' is a sequential′ equilibrium in that game. However, if instead of looking at the information sets $\{x_1\}$ and $\{x_2\}$ individually, we consider both of them, and require that a strategy do optimally conditional on reaching $\{x_1, x_2\}$ (the union of these information sets), then the only strategy that does so is b; b' does not meet this requirement.

To make this precise, we need to generalize the definition of a belief system. Recall that a belief system μ for a game Γ associates with each information set X in Γ a probability μ_X on the histories in X. Such a belief system does not suffice if we need to compute whether b'_i does better conditional on reaching a set $\mathcal{X}$ of information sets. A *generalized belief system* μ for a game Γ associates with each (non-empty) set $\mathcal{X}$ of information sets in Γ a probability $\mu_{\mathcal{X}}$ on the histories in the union of the information sets $X_i \in \mathcal{X}$. As before, we interpret $\mu_{\mathcal{X}}(x)$ as the probability of reaching x conditioned on reaching $\mathcal{X}$ and require that $\sum_{x \in \hat{\mathcal{X}}} \mu_{\mathcal{X}}(x) = 1$, where $\hat{\mathcal{X}}$ denotes the upper frontier of histories in $\mathcal{X}$, that is, all the nodes $x \in \cup\mathcal{X}$ such that there is no node $x' \in X$ with $x' \prec x$. We can now define expected utility in exactly the same way as before.

As before, we say that a generalized belief system μ is *compatible* with a strategy b if, for all information sets $\mathcal{X}$ of informations sets such that $\pi_b(\mathcal{X}) > 0$, we have $\mu_{\mathcal{X}}(x) = \pi_b(x \mid \mathcal{X})$. The analogue of Proposition 3.3 holds:

Proposition 4.7: *If a set $\mathcal{X}$ of information sets is reached by strategy profile b with positive probability, and μ is compatible with b, then*

$$EU_i(b \mid \mathcal{X}) = EU_i(b \mid \mathcal{X}, \mu).$$

Proof: The proof is identical to the proof of Proposition 3.3. ∎

All the other notions we introduced generalize in a straightforward way. A generalized belief system μ is *consistent* with a strategy b if there exists a sequence of completely mixed strategy profiles $b^1, b^2, \ldots$ converging to b such that $\mu_{\mathcal{X}}(x)$ is $\lim_{n\to\infty} \pi_{b^n}(x \mid \mathcal{X})$. We say that $\mathcal{X}$ *precedes* X', written $\mathcal{X} \preceq X'$, iff for all $x' \in X'$ there exists some $x \in \cup\mathcal{X}$ such that x precedes x' on the game tree; that is, we are defining $\preceq$ exactly as before, identifying the set $\mathcal{X}$ with the union of the information sets it contains. As before, we define $[b, \mathcal{X}, b']$ to be the strategy according to which i plays b' at every information set X' such that $\mathcal{X} \preceq X'$, and otherwise plays b. Let $(b, \mathcal{X}, b')$ denote the "strategy" of using b until $\mathcal{X}$ is reached and then switching to b'. The analogue of Proposition 4.1 holds: $(b, \mathcal{X}, b')$ is a strategy in game Γ iff $(b, \mathcal{X}, b') = [b, \mathcal{X}, b']$.

We now have the following generalization of Proposition 4.3.

Proposition 4.8: *If b^* is a perfect equilibrium in game Γ, then there exists a generalized belief system μ consistent with b^* such that, for all players i, all non-empty sets $\mathcal{X}$ of information sets for player i, and all behavioral strategies b' for player i at X, we have*

$$EU_i(b \mid \mathcal{X}, \mu) \geq EU_i(([b_i, \mathcal{X}, b'], b_{-i}) \mid \mathcal{X}, \mu).$$

Proof: The proof is identical to that of Proposition 4.3, except that we use Proposition 4.7 instead of Proposition 3.3. ∎

We can now formally define sequential equilibrium.

Definition 4.9: A pair (b, μ) consisting of a strategy profile b and a generalized belief system μ is called a *generalized belief assessment*. A generalized belief assessment (b, μ) is a *sequential equilibrium* in a game Γ if μ is consistent with b and for all players i, all non-empty sets $\mathcal{X}$ of information sets for player i, and all behavioral strategies b' for player i, we have

$$EU_i(b \mid \mathcal{X}, \mu) \geq EU_i(([b_i, \mathcal{X}, b'], b_{-i}) \mid \mathcal{X}, \mu).$$

∎

It is immediate from Proposition 4.8 that every perfect equilibrium is a sequential equilibrium. Thus, every game has a sequential equilibrium.

Theorem 4.10: *A sequential equilibrium exists in all finite games.*

It is immediate from the definitions that every sequential equilibrium is a sequential$'$ equilibrium. Furthermore, as the definition of sequential equilibrium considers changes at all sets of information sets, and in particular, the set consisting of all information sets, it follows that every sequential equilibrium is a Nash equilibrium. (Recall that this was not the case for sequential$'$ equilibrium.) Finally, we note that if (b, μ) is a sequential$'$ equilibrium of a game of perfect recall Γ (so that it is a sequential equilibrium in the sense of Kreps and Wilson (1982)), then there exists a generalized belief system μ' such that (b, μ') is a sequential equilibrium in Γ in the sense that we have just defined: Consider the sequence of strategy profiles $b^1, b^2, \ldots$ that define μ; this sequence also determines a generalized belief system μ'. We claim that (b, μ') is a sequential equilibrium in our sense. If not, there exists some player i, a set $\mathcal{X}$ of information sets for i, and a behavioral strategy b'_i, such that conditional on reaching $\mathcal{X}$, i prefers using b', given belief assessment μ'. This implies that there exists an information set $X \in \mathcal{X}$ such that i also prefers switching to b' at X, given belief assessment μ'. But μ and μ' assign the same beliefs to the information set X (since they are defined by the same sequence of strategy profiles), which means that i also prefers switching to b' at X, given belief assessment μ, so (b, μ) cannot be a sequential$'$ equilibrium. We conclude that in games of perfect recall, every sequential$'$ equilibrium is also a sequential equilibrium.

As we noted earlier, this is no longer true in games of imperfect recall—in the game in match nature game, b' is a sequential$'$ equilibrium, but is not a sequential equilibrium. The argument above fails because for games of imperfect recall, (b_i, X', b') (i.e., switching from b_i to b' at information set X') might not be a valid strategy even if $(b_i, \mathcal{X}, b')$ is; this cannot happen in games of perfect recall.

Summarizing this discussion, we have the following result.

Theorem 4.11: *In all finite games,*

(a) every sequential equilibrium is a Nash equilibrium;

(b) in games of perfect recall, a strategy profile b is a sequential equilibrium iff it is a sequential$'$ equilibrium.

However, there exist games of imperfect recall where a sequential$'$ equilibrium is not a sequential equilibrium.

5 Discussion

Selten (1975) says that "game theory is concerned with the behavior of absolutely rational decision makers whose capabilities of reasoning and remembering are unlimited, a game . . . must have perfect recall." We disagree. We believe that game theory ought to be concerned with decision makers that may not be absolutely rational and, more importantly for the present paper, players that do not have unlimited capabilities of reasoning and remembering. This is particularly critical when applying game-theoretic ideas in computer science. Although there has been relatively little work on imperfect recall in game theory, imperfect recall arises in arguably the standard setting in computer science. Computers have bounded recall (as do people)! Defining appropriate solution concepts in the presence of imperfect recall is notoriously tricky, as the work of PR shows. What seems optimal to an agent *ex ante* may no longer seem optimal once the agent reaches a particular information set.

In this paper, we have defined *ex ante* notions of sequential equilibrium and perfect equilibrium. We have also pointed out the subtleties in doing so. We did so in the standard game-theoretic model of extensive-form games with information sets. A case can be made that the problems that arise in defining sequential equilibrium stem in part from the use of the standard framework, which models agents' information using information sets (and then requires that agents act the same way at all nodes in an information set). This does not allow us to take into account, for exampe, whether or not an agent knows his strategy. Halpern (1997) shows that many of the problems pointed out by PR can be dealt with using a more "fine-grained" model, the so-called *runs-and-systems* framework (Fagin et al. 1995), where agents have local states that characterize their information. The local state can, for example, include the agents' strategy (and modifications to it). It would be interesting to explore how the ideas of this paper play out in the runs-and-systems framework. We have taken preliminary steps to doing this in a computational setting (Halpern and Pass 2013), but clearly more work needs to be done to understand what the "right" solution concepts are in a computational setting.

References

Aumann, R. J.; Hart, S.; and Perry, M. 1997. The absent-minded driver. *Games and Economic Behavior* 20:102–116.

Fagin, R.; Halpern, J. Y.; Moses, Y.; and Vardi, M. Y. 1995. *Reasoning About Knowledge*. Cambridge, Mass.: MIT Press. A slightly revised paperback version was published in 2003.

Gilboa, I. 1997. A comment on the absent-minded driver paradox. *Games and Economic Behavior* 20(1):25–30.

Grove, A. J., and Halpern, J. Y. 1997. On the expected value of games with absentmindedness. *Games and Economic Behavior* 20:51–65.

Halpern, J. Y., and Pass, R. 2013. Sequential equilibrium in computational games. In *Proc. Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI '13)*, 171–176.

Halpern, J. Y. 1997. On ambiguities in the interpretation of game trees. *Games and Economic Behavior* 20:66–96.

Hendon, E.; Jacobsen, H. J.; and Sloth, B. 1996. The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior* 12(2):274–282.

Isbell, J. 1957. Finitary games. In Dresher, M., and Kuhn, H. W., eds., *Contributions to the Theory of Games III*. Princeton, N.J.: Princeton University Press. 79–96.

Kaneko, M., and Kline, J. J. 1995. Behavior strategies, mixed strategies, and perfect recall. *International Journal of Game Theory* 24:127–145.

Kline, J. J. 2005. Imperfect recall and the relationships between solution concepts in extensive games. *Economic Theory* 25:703–710.

Kreps, D. M., and Wilson, R. B. 1982. Sequential equilibria. *Econometrica* 50:863–894.

Lipman, B. L. 1997. More absentmindedness. *Games and Economic Behavior* 20(1):97–101.

Osborne, M. J., and Rubinstein, A. 1994. *A Course in Game Theory*. Cambridge, Mass.: MIT Press.

Piccione, M., and Rubinstein, A. 1997a. The absent-minded driver's paradox: synthesis and responses. *Games and Economic Behavior* 20(1):121–130.

Piccione, M., and Rubinstein, A. 1997b. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20(1):3–24.

Rockafellar, R. T. 1970. *Convex Analysis*. Princeton, N. J.: Princeton University Press.

Rubinstein, A. 1986. Finite automata play the repeated prisoner's dilemma. *Journal of Economic Theory* 39:83–96.

Selten, R. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4:25–55.

Wichardt, P. C. 2008. Existence of Nash equilibria in finite extensive form games with imperfect recall: A counterexample. *Games and Economic Behavior* 63(1):366–369.