

Sequential Equilibrium in Games of Imperfect Recall*

Joseph Y. Halpern[†]
Cornell University
halpern@cs.cornell.edu

Rafael Pass[‡]
Cornell University
rafael@cs.cornell.edu

March 7, 2018

Abstract

Sequential equilibrium is defined in games of imperfect recall. Subtleties regarding the definition are discussed.

JEL Classification numbers: C70, D81.

1 Introduction

Sequential equilibrium [Kreps and Wilson 1982] is perhaps the most common solution concepts used in extensive-form games. It is trying to capture the intuition that agents play optimally, not just on the equilibrium path, but also off the equilibrium path. Unfortunately, sequential equilibrium has 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

*A shorter preliminary version of this paper appeared in *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning (KR 2016)*, 2016, pp. 278-287.

[†]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.

are even modestly longer than the standard two- or three-move games considered in most game theory papers. Nevertheless, the intuition that underlies these solution concepts, 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 single-player game described in Figure 1, which we call the “match-nature” game, where nature makes an initial move, either left or right, and then the agent moves. The agent can stop the game (by playing S) immediately after nature’s move, or can continue. If the agent continues, it does best by matching nature’s initial move (hence the name). However, although the agent knew nature’s move initially, he forgets it if he continues. Thus, this is a game of imperfect recall.

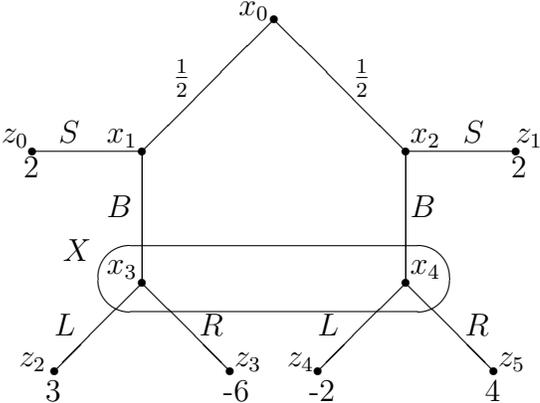


Figure 1: Subtleties with imperfect recall, illustrated by the match-nature game.

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 (playing S at x_1 , B at x_2 , and then R at the information set X) and then switches to b' (playing B at x_1 and L at X) 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.

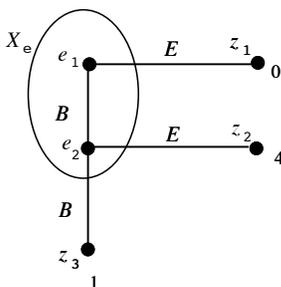


Figure 2: The absentminded driver game.

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 behavioral 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

¹As 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.

whether he should switch to a new strategy, where he exits with probability p . Another straightforward calculation shows that his expected payoff is then

$$\alpha((1-p)^2 + 4p(1-p)) + (1-\alpha)((1-p) + 4p) = 1 + (3-\alpha)p - 3\alpha p^2. \quad (1)$$

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. ■

Our goal in this paper is to define sequential equilibrium for games of imperfect recall. As these examples show, such a definition requires a clear interpretation of the meaning of information sets and the restrictions they impose on the knowledge and strategies of players. But there are further subtleties in games of imperfect recall. Intuitively, in a sequential equilibrium, an agent feels that she is playing optimally even off the equilibrium path. To determine optimal play, in each of her information sets, an agent must have beliefs regarding how likely she is to be at each node. One issue that must be addressed is how these beliefs should be determined.

According to the technique used by Selten [1975], 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.)

While the issue of how the agent should ascribe beliefs has been considered at length in the literature, an issue that has received less attention is what strategies an agent can deviate to at an information set. 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 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. If the strategy profile b is a sequential equilibrium and X is an information set for

i , then we want i 's utility with b to be at least as high as i 's utility with $([b_i, X, b'_i], b_{-i})$, where, intuitively, $[b_i, X, b'_i]$ is the strategy where the agent plays b up to X and then continues with b' at and after X (and, as usual, the other agents continue to play b_{-i} , their component of b). But what does “after X ” mean? If X_1 and X_2 are information sets in a game of perfect recall and there is a path p that includes a node $x_1 \in X_1$ and a node $x_2 \in X_2$ such that x_1 precedes x_2 on p , then for every path p' that includes a node $x'_2 \in X_2$, there is a node $x'_1 \in X_1$ that precedes x'_2 on p' . This allows us to put a partial order on information sets in games of perfect recall such that making a change at and after an information set leads to a well-defined strategy.

However, this is in general not the case in games of imperfect recall. Consider the match-nature game. Although there is a path in the game tree where x_1 precedes a node in X (namely, x_2), there is a path in the tree that includes a node in X (namely, x_3) that does not include x_1 . This is precisely why making a change at and after x_1 “breaks” the information structure of the game.

Our *ex ante* notion of sequential equilibrium does not allow changes that break the information structure of the game. That is, at an information set X , we only allow changes to strategies $[b, X, b']$ that are compatible with the information structure of the game. This approach also leads to a natural way of defining the agent's beliefs at an information set.

Unfortunately, this approach 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 Section 4), 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. (We discuss this issue in more detail in Section 4.5.) Moreover, as we argue briefly in Section 5, 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].

There is a final issue we must confront when deal with imperfect recall in full generality. As is well known, in games of perfect recall, behavioral strategies and mixed strategies are *outcome-equivalent* [Kuhn 1953] (see Section 2.3 for formal definitions), but once we move to games of imperfect recall, they are incomparable in expressive power [Wichardt 2008]. To show the existence of Nash equilibrium, we must move to *mixed behavioral strategies*, that is, distributions over behavioral strategies. Thus, we work in this paper with mixed behavioral strategies, which leads to a few additional technical complications.

The rest of this paper is organized as follows. In Section 2, we expand briefly on a number of the issues touched on above, such as belief ascription and mixed behavioral strategies; these preliminaries will be necessary for understanding our formal definition

of (*ex ante*) sequential equilibrium. We provide that formal definition in Section 4. It turns out that defining *quasi-perfect equilibrium* [van Damme 1984] in games of imperfect recall is a useful tool in defining sequential equilibrium, so we start with that definition. We then give a preliminary definition of an *ex ante* sequential equilibrium, and show that it exists in all games (including games of imperfect recall). Of course, in games of perfect recall, the notion agrees with the standard definition sequential equilibrium. But a closer examination of our preliminary definition reveals a further subtlety: not all *ex ante* sequential equilibria are Nash equilibria under our initial definition. As we show in Section 4.4, we need to allow simultaneous changes to a strategy at and below more than one information set at a time to get what is arguably a more appropriate notion of sequential equilibrium. We then discuss interim sequential equilibrium in Section 4.5. We conclude with some discussion in Section 5.

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

We assume that the reader is familiar with the standard definition of extensive-form games and perfect recall in such games (see, for example, [Osborne and Rubinstein 1994] for a formal treatment). Recall that 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. 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 histories h and h' in the same information set and h' is a prefix of h .² 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 assumption tends to be explicit in epistemic analyses of game theory; it arises in much of the discussion of imperfect recall as well. For simplicity,

²Following Osborne and Rubinstein [1994], formally, we speak in terms of histories, which are sequences of actions, rather than nodes on the game tree. Each history leads to a unique node in the game tree. If history h' is a prefix of h , then the node to which h' leads is an ancestor in the game tree of the node to which h leads. We occasionally talk about a node being *on a history*; this just means that the history goes through that node in the game tree.

consider one-player games, that is, decision problems, with perfect recall. Then it could be argued that players do not really need to know their strategies. After all, a rational player could just compute at each information set X what the *ex ante* optimal strategy is, and then play the move that it recommends at X . If the optimal move is not unique, there is no problem—any choice of optimal move will do.

Things change when we move to games of imperfect recall. 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 . It could be argued that, as long as the agent remembers the structure of the game, then, just as in the case of perfect recall, he can recompute the *ex ante* optimal strategy at each information set X and play the move that it recommends at X . However, we now run into a problem if the optimal strategy is not unique. With imperfect recall, if there are ties, it may well matter which choice the agent makes. For example, suppose that 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 !

For an *ex ante* notion of sequential equilibrium, it seems arguably more reasonable to assume that the agent commits initially to playing the equilibrium strategy (and will know this strategy at later nodes in the game tree). But we stress that this assumption is 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.

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] gives an example of a game with imperfect recall with no Nash equilibrium in behavioral strategies.

Thus, to deal with games of imperfect recall, in general, we need to allow *behavioral strategy mixtures* [Isbell 1957], which are distributions over behavioral strategies.³ As Kaneko and Kline [1995] 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. We try to consistently use the notation c to denote behavioral strategy mixtures and b to denote behavioral strategies that are not mixtures.

Players do not have to mix over too many behavioral strategies when employing a behavioral strategy mixture. Alpern [1988] proves the following result:

Proposition 2.1: *If Γ is a finite game, there is a constant D_Γ that depends only on Γ such that each behavioral strategy mixture is outcome-equivalent to a behavioral strategy mixture that mixes over at most D_Γ strategies.*⁴

³Isbell [1957] actually calls them mixed strategies. Selten [1975] called them *behavior strategy mixtures*. We write “behavioral” for consistency with our terminology elsewhere.

⁴Alpern [1988] states a weaker result: that the components of a Nash equilibrium in a two-person zero-sum game are equivalent to a mixture of a bounded number of behavioral strategies. However, his proof shows the stronger result that we have claimed. We thank John Hillas for pointing us to Alpern’s

A consequence of Proposition 2.1 is that in a finite game Γ where player i can play d_i possible strategies, we can identify a mixed behavioral strategy for player i with an element of $([0, 1] \times R^{d_i})^{D_\Gamma}$ —each mixed behavioral strategy can be viewed as a tuple of the form $((a_1, b_1), \dots, (a_{D_\Gamma}, b_{D_\Gamma}))$, where $a_1, \dots, a_{D_\Gamma} \in [0, 1]$, $\sum a_i = 1$, and b_j is a mixed behavioral strategy for player i , and thus in R^{d_i} , for $j = 1, \dots, D_G$. Since it is well known that the convex hull 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.

Solution concepts such as Nash equilibrium and sequential equilibrium 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. This will also be the case for the solution concepts that we define. Thus, in light of Proposition 2.1, when considering a game Γ , we consider only behavioral strategy mixtures whose support contains at most D_Γ behavioral strategies.

Sequential equilibrium is usually defined in terms of behavioral strategies, not mixed 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* behavioral strategy choice.

2.4 Expected Utility of Strategies

Every behavioral strategy mixture profile c induces a probability measure π_c on leaves. 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 $\pi_c(x)$ to be the probability of reaching a leaf that comes after x when using strategy c .

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 $EU_i(c)$ to be $\sum_{z \in Z} \pi_c(z)u_i(z)$. If Y is a subset of leaves such that $\pi_c(Y) > 0$, then computing the expected utility of c for player i conditional on Y is equally straightforward. It is simply

$$EU_i(c | Y) = \sum_{z \in Y} \pi_c(z | Y)u_i(z);$$

result.

that is, the expected utility of c for player i conditional on Y is just the sum, over all terminal histories z in Y , of the probability of c conditional on Y times the utility of z for i .

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 that do not exhibit absentmindedness, 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 c with positive probability and $\mu_X(x) = \pi_c(x | X)$, then $\sum_{x \in \hat{X}} \mu_X(x) = 1$.

Proof: By definition, $\sum_{x \in \hat{X}} \mu_X(x) = \sum_{x \in \hat{X}} \pi_c(x | X) = \pi_c(\hat{X} | X) = 1$. ■

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

Our definition of μ_X^c 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 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^c 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 c' .

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 c , we have $\sum_{z \in Z} \mu_X^c(z) = 1$.

Proof: By definition,

$$\begin{aligned} \sum_{z \in Z} \mu_X^c(z) &= \sum_{z \in Z} \mu_X(x_z) \pi_c(z | x_z) \\ &= \sum_{z \in Z} \sum_{x \in \hat{X}} \mu_X(x) \pi_c(z | x) \\ &= \sum_{x \in \hat{X}} \mu_X(x) \sum_{z \in Z} \pi_c(z | x) \\ &= \sum_{x \in \hat{X}} \mu_X(x) \\ &= 1. \end{aligned}$$

■

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

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

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

$$EU_i(c | X) = EU_i(c | X, \mu).$$

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

$$\begin{aligned} &EU_i(c | X, \mu) \\ &= \sum_{z \in Z} \mu_X^c(z) u_i(z) \\ &= \sum_{z \in Z_X} \mu_X(x_z) \pi_c(z | x_z) u_i(z) \\ &= \sum_{z \in Z_X} \pi_c(x_z | X) \pi_c(z | x_z) u_i(z) \\ &= \sum_{z \in Z_X} \pi_c(z | X) u_i(z) \\ &= EU_i(c | X). \end{aligned}$$

■

4 Quasi-Perfect, Perfect, and Sequential Equilibrium

4.1 Perturbed Games

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 sequential 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 a that can be performed at X , a probability $\eta_a \geq 0$ 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 $\sum_{a \in A(X)} \eta_a < 1$. We call η a *perturbation* of Γ . We think of η_a as the probability of a "tremble"; since we view trembles as unlikely, we are most interested in the case that η_a is small but positive.

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 $a \in A(X)$, $b(X)$ assigns probability greater than or equal to η_a to a . A behavioral strategy mixture c is acceptable in (Γ, η) if each behavioral strategy in its support is acceptable in (Γ, η) . The proof of Proposition 2.1 given by Alpern [1988] shows a slightly stronger result:

Proposition 4.1: *If Γ is a finite game, there is a constant D_Γ that depends only on Γ such that, for all perturbations η , every behavioral strategy mixture acceptable in (Γ, η) is outcome-equivalent to a behavioral strategy mixture acceptable in (Γ, η) that mixes over at most D_Γ strategies.*

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 c is an acceptable strategy profile in a perturbed game (Γ, η) where $\eta > 0$ (i.e., $\eta_a > 0$ for all actions a), then $\pi_c(X) > 0$ for all information sets X , and if $\eta = 0$ then the acceptable strategy profiles in (Γ, η) are just the strategy profiles in Γ .

4.2 Best Responses at Information Sets

There are a number of ways to capture the intuition that a behavioral 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 $b_i[X/a]$ denote the behavioral strategy that is identical to b_i except that, at information set X , action a 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 behavioral 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 $X_2 \prec' X_1$.⁵

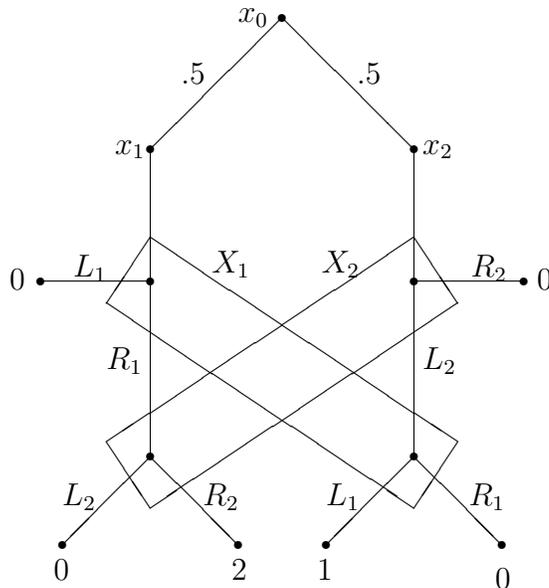


Figure 3: A game where \prec' is not a partial order.

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

⁵Okada [1987] defined a notion called by Kline [2005] *occurrence memory*; a partition satisfies occurrence memory if, for all information sets X and X' , if $X \prec' X'$, then $X \prec X'$. Thus, if a partition satisfies occurrence memory, \prec is equivalent to \prec' . Kline interprets occurrence memory as saying that an agent recalls what he learned (but perhaps not what he did).

behavioral strategy. 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 . We do not define $[c, X, c']$ if either c or c' is a mixed behavioral strategy (as opposed to just a behavioral strategy).

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 4.2: *If b and b' are behavioral strategies in game Γ , then (b, X, b') is a behavioral strategy iff $(b, X, b') = [b, X, b']$.*

Proof: 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. ■

4.3 Defining Quasi-Perfect and Sequential Equilibrium

Before defining sequential equilibrium in games of imperfect recall, it is useful to define quasi-perfect equilibrium.

4.3.1 Quasi-perfect equilibrium

To define quasi-perfect equilibrium, we can use literally the same definition as van Damme [1984], except that we use behavioral strategy mixtures rather than behavioral strategies. The behavioral strategy mixture profile c^* is a *quasi-perfect equilibrium* in Γ if there exists a sequence $(\Gamma, \eta_1), (\Gamma, \eta_2), \dots$ of perturbed games and a sequence of behavioral strategy mixture profiles c^1, c^2, \dots such that (1) $\eta_k \rightarrow \vec{0}$; (2) c^k is a Nash equilibrium of (Γ, η_k) ;

and (3) $c^k \rightarrow c^*$.⁶ Van Damme [1984] shows that a quasi-perfect equilibrium always exists in games with perfect recall. Essentially the same proof shows that it exists even in games with imperfect recall.

Theorem 4.3: *A quasi-perfect equilibrium exists in all finite games.*

Proof: Consider any sequence $(\Gamma, \eta_1), (\Gamma, \eta_2), \dots$ 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 c^k in behavioral strategy mixtures. Here we are using the fact that, by Proposition 4.1, the set of behavioral strategy mixtures acceptable in (Γ, η_k) is compact. Moreover, there is a single compact space such C that all the behavioral strategy mixtures are (outcome-equivalent to a behavioral strategy mixture) in C . By a standard compactness argument, the sequence c^1, c^2, \dots has a convergent subsequence. Suppose that this subsequence converges to c^* . Clearly c^* is a quasi-perfect equilibrium. ■

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 quasi-perfect equilibrium is a Nash equilibrium, it follows from Wichardt's [2008] example that a quasi-perfect equilibrium may not exist in a game 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 quasi-perfect equilibrium c^* of Γ is also a Nash equilibrium of Γ . Thus, each strategy c_i^* is a best response to c_{-i}^* *ex ante*. However, we also want c_i^* to be a best response to c_{-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 $a \in A(X)$, b_i assigns positive probability to playing a . 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 behavioral strategy mixture profile c if there exists a sequence of completely mixed strategy profiles c^1, c^2, \dots converging to c such that $\mu_X(x)$ is $\lim_{n \rightarrow \infty} \pi_{c^n}(x | X)$. Note that if μ is consistent with c , then it is compatible with c .

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

Proposition 4.4: *If the behavioral strategy mixture profile c is a quasi-perfect equilibrium in game Γ , then for all behavioral strategies b_i in the support of c_i , there exists a belief*

⁶We 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.

system μ_{b_i} consistent with (b_i, c_{-i}) such that, for all players i , all information sets X for player i , and all behavioral strategies b' for player i , we have

$$EU_i((b_i, c_{-i}) | X, \mu_{b_i}) \geq EU_i([b_i, X, b'], c_{-i} | X, \mu_{b_i}).^7$$

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

We claim that the result holds with respect to μ_{b_i} . For suppose not. Then there exists a player i , information set X , behavioral strategy b' for player i , and $\epsilon > 0$ such that $EU_i((b_i, c_{-i}) | X, \mu_{b_i}) + \epsilon < EU_i([b_i, X, b'], c_{-i} | X, \mu_{b_i})$. It follows from Proposition 3.3 that $EU_i((b_i^k, c_{-i}^k) | X) \rightarrow EU_i((b_i, c_{-i}) | X, \mu_{b_i})$ and $EU_i([b_i^k, X, b'], c_{-i}^k | X) \rightarrow EU_i([b_i, X, b'], c_{-i} | X, \mu_{b_i})$. Since $(b_i^k, c_{-i}^k) \rightarrow (b_i, c_{-i})$ and $\eta_k \rightarrow \vec{0}$, there exists some strategy b'' and $k > 0$ such that b'' is acceptable for $(\Gamma, \eta_{k'})$ for all $k' > k$, and $EU_i((b_i^{k'}, c_{-i}^{k'}) | X) + \epsilon/2 < EU_i([b_i^{k'}, X, b''], c_{-i}^{k'} | X)$. But this contradicts the assumption that $c^{k'}$ is a Nash equilibrium of $(\Gamma, \eta_{k'})$ (since this assumption also implies that $EU_i(c^{k'}) = EU_i((b_i^{k'}, c_{-i}^{k'}))$).⁸ ■

We are implicitly identifying “ b_i is a best response for i at information set X ” with “ $EU_i((b_i, c_{-i}) | X, \mu_{b_i}) \geq EU_i([b_i, X, b'], c_{-i} | X, \mu_{b_i})$ ” for all behavioral strategies b' for player i . How reasonable is it to consider $[b_i, X, b']$ here? 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 recall. For example, in the match-nature game, the player's action at x_3 can clearly affect his payoff conditional on reaching x_1 , but the information set X that contains x_3 does not come after $\{x_1\}$, so we do not allow changes at x_3 in considering best responses at x_1 . While making a change at x_3 makes things better at x_1 , it would make things worse

⁷In an earlier draft of this paper, we consider a belief assessment μ consistent with c rather than a belief assessment μ_{b_i} consistent with (b_i, c_{-i}) . The latter requirement models the fact that, once player i chooses a behavioral strategy at the beginning of the game, he remembers his choice throughout the game (and, in particular, knows it at information set X). We thank John Hillas for emphasizing the need to consider μ_{b_i} .

⁸Note 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).

at x_2 , a node that is not after x_1 . Given our *ex ante* viewpoint, this is clearly a relevant consideration. What we are really requiring is that b_i is a best response for i to c_{-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. The strategy b is a best response at x_1 among strategies that do not affect i 's utility at nodes that do not come after $\{x_1\}$; although b' gives the agent a higher utility than b conditional on reaching x_1 , it affects i 's utility at x_3 . We return to this point in Section 4.4.

We can now define a notion of sequential equilibrium essentially as Kreps and Wilson [1982] did. The two definitions are literally identical for profiles of behavioral strategies, and easily seen to be equivalent to Kreps and Wilson's definition in games of perfect recall, but we must generalize it appropriately for profiles of behavioral strategy mixtures in games of imperfect recall. Moreover, this notion turns out to not imply Nash equilibrium in games of imperfect recall, 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 4.5: A *belief assessment family* is pair $(c, \{\mu_{b_i} : b_i \text{ is in the support of } c_i\})$ consisting of a behavioral strategy mixture profile c and a family of belief systems μ_{b_i} , one for each player i and behavioral strategy b_i in the support of c_i . A belief assessment family $(c, \{\mu_{b_i} : b_i \text{ is in the support of } c_i\})$ is a *sequential' equilibrium* in a game Γ if μ_{b_i} is consistent with (b_i, c_{-i}) and, for all players i , all information sets X for player i , all behavioral strategies b_i in the support of c_i , and all behavioral strategies b' for player i at X , we have

$$EU_i((b_i, c_{-i}) \mid X, \mu_{b_i}) \geq EU_i((b_i, X, b'), c_{-i} \mid X, \mu_{b_i}).$$

■

In games of perfect recall, it is easy to see that we can take $\mu_{b_i} = \mu_{b'_i}$; agent i 's beliefs at an information set X just depend on the sequence of actions played by agent i in reaching information set X , and (with perfect recall) this sequence must be the same no matter which pure (or behavioral) strategy the agent chose to play initially. Thus, in games of perfect recall, a belief assessment family can be identified with the Kreps and Wilson notion of a belief assessment. However, there are examples showing that this is not the case in games of imperfect recall (see [Hillas and Kvasov 2017]).

It is immediate from Proposition 4.4 that every quasi-perfect equilibrium is a sequential' equilibrium. Thus, a sequential' equilibrium exists for every game.

Theorem 4.6: *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 .

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 4.7: *In all finite games,*

- (a) *every quasi-perfect equilibrium is a Nash equilibrium;*
- (b) *there exist games where a sequential' equilibrium is not a Nash equilibrium;*
- (c) *in games where all players have an initial information set, every sequential' equilibrium is a Nash equilibrium.*

It is 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. Consider the modification of the match-nature game with an initial node x_{-1} . At x_{-1} , the agent can only play down, leading to x_0 . This game has only one Nash equilibrium: playing down then b . By Theorem 4.7(b), this is also the only sequential' equilibrium of the game. However, replacing b by b' gives a strategy that satisfies the one-step deviation principle but is not a sequential' equilibrium.

We remark that we could also define a notion of *perfect equilibrium* [Selten 1975] in games of imperfect recall. However, as van Damme [1984] observes, perfect equilibrium in extensive-form games can be understood in terms of local best responses. Specifically, he shows that a behavioral strategy profile b is a perfect equilibrium in a game of perfect recall iff there exists a sequence $b^k \rightarrow b$ of completely mixed strategies such that, for every agent i and every information set X for agent i , the strategy $b_i^k[X/b_i(X)]$, which is just like b_i^k except that, at information set X , player i plays the distribution $b_i(X)$, is a

best response to b_{-i}^k conditional on reaching X . The fact that perfect equilibrium implies Nash equilibrium in games of perfect recall then follows from the one-step deviation principle. But since this principle does not hold in games of imperfect recall, we should not expect the obvious extension of the notion of perfect equilibrium to imply Nash equilibrium in games of imperfect recall. And, indeed, the strategy b' in the match-nature game satisfies van Damme's characterization of perfect equilibrium, but is not a Nash equilibrium. Quasi-perfect equilibrium allows an agent to make coordinated deviations at several information sets. It turns out that considering several information sets is the key to getting a reasonable notion of sequential equilibrium, as we show in the next section.

4.4 Sequential Equilibrium

Our *ex ante* notion of sequential' equilibrium does not allow the agent to switch to a strategy b' at an information set X if doing so would affect his utility at nodes that do not come after X . Such a switch would result in a "strategy" incompatible with the information structure of the game. While we do not want to allow an agent to use a strategy incompatible with the information structure of the game, an agent may be able to switch to a strategy compatible with the game's information structure by considering changes not just at and below one information set at a time, but by considering changes at several information sets simultaneously. Allowing changes at several information sets has no impact on the notion of sequential' equilibrium in games of perfect recall, but it can have a significant impact in games of imperfect recall. Moreover, such considerations seem to us in the spirit of our *ex ante* viewpoint.

These points can perhaps be best clarified by looking at an example. Consider the match-nature game again. As we observed, the strategy b' is a sequential' equilibrium in that game. However, suppose that rather than looking at the information sets $\{x_1\}$ and $\{x_2\}$ individually, we allow changes at both of them; that is, we allow an agent that is using strategy b_1 to switch to a strategy b_2 at and below $\{x_1, x_2\}$ (the union of the two information sets) provided that the switch does not affect the agent's utility at nodes that do not come after $\{x_1, x_2\}$. It is easy to see that the only strategy that satisfies this stronger requirement is b ; b' does not satisfy it.

As we now show, allowing simultaneous changes at sets of information sets allows us to define a notion of sequential equilibrium that generalizes Nash equilibrium not just in the match-nature game, but in all games. As a first step to formalizing these ideas, 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 $\mathcal{X} = \{X_1, \dots, X_k\}$ 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 profile c if, for all information sets \mathcal{X} of information sets such that $\pi_c(\mathcal{X}) > 0$, we have $\mu_{\mathcal{X}}(x) = \pi_c(x | \mathcal{X})$. The analogue of Proposition 3.3 holds:

Proposition 4.8: *If a set \mathcal{X} of information sets is reached by strategy profile c with positive probability, and μ is compatible with c , then*

$$EU_i(c | \mathcal{X}) = EU_i(c | \mathcal{X}, \mu).$$

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 behavioral strategy mixture profile c if there exists a sequence of completely mixed strategy profiles c^1, c^2, \dots converging to c such that $\mu_{\mathcal{X}}(x)$ is $\lim_{n \rightarrow \infty} \pi_{c^n}(x | \mathcal{X})$. 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 behavioral 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 4.2 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 4.4.

Proposition 4.9: *If the behavioral strategy mixture profile c is a quasi-perfect equilibrium in game Γ , then for all behavioral strategies b_i in the support of c_i , there exists a generalized belief system μ_{b_i} consistent with (b_i, c_{-i}) 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*

$$EU_i((b_i, c_{-i}) | \mathcal{X}, \mu_{b_i}) \geq EU_i([b_i, \mathcal{X}, b'], c_{-i} | \mathcal{X}, \mu_{b_i}).$$

Proof: The proof is identical to that of Proposition 4.4, except that we use Proposition 4.8 instead of Proposition 3.3. ■

We can now formally define sequential equilibrium.

Definition 4.10: A pair $(c, \{\mu_{b_i} : b_i \text{ is in the support of } c_i\})$ consisting of a behavioral strategy mixture profile c and a family of belief systems μ_{b_i} , one for each player i and behavioral strategy b_i in the support of c_i , is called a *generalized belief assessment family*. A generalized belief assessment family $(c, \{\mu_{b_i} : b_i \text{ is in the support of } c_i\})$ is a *sequential equilibrium* in a game Γ if μ_{b_i} is consistent with (b_i, c_{-i}) 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

$$EU_i((b_i, c_{-i}) | \mathcal{X}, \mu_{b_i}) \geq EU_i([b_i, \mathcal{X}, b'], c_{-i} | \mathcal{X}, \mu_{b_i}).$$

■

It is immediate from Proposition 4.9 that every quasi-perfect equilibrium is a sequential equilibrium. Thus, every game has a sequential equilibrium.

Theorem 4.11: *A sequential equilibrium exists in all finite games.*

Clearly 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, suppose that $(c, \{\mu_{b_i} : b_i \text{ is in the support of } c_i\})$ is a sequential' equilibrium of a game of perfect recall Γ . Then, as we observed earlier, we can take $\mu_{b_i} = \mu_{b'_i}$ for all behavioral strategies b_i and b'_i in the support of c_i , so we can identify the set $\{\mu_{b_i} : b_i \text{ is in the support of } c_i\}$ with a single belief system μ , and (c, μ) is a sequential equilibrium in the sense of Kreps and Wilson [1982]. In that case, there exists a generalized belief system μ' such that (c, μ') is a sequential equilibrium in Γ in the sense that we have just defined (identifying μ' with the set $\{\mu'_{b_i} : b_i \text{ is in the support of } c_i\}$, where $\mu'_{b_i} = \mu'$ for all b_i). Consider the sequence of strategy profiles c^1, c^2, \dots that define μ ; this sequence also determines a generalized belief system μ' . We claim that (c, μ') is a sequential equilibrium in our sense. If not, there exists some player i , a set \mathcal{X} of information sets for i , a strategy b_i in the support of c , and a behavioral strategy b'_i , such that conditional on reaching \mathcal{X} , i prefers using b' , given belief assessment belief assessment μ' . This implies that there exists an information set $X \in \mathcal{X}$ such that i also prefers switching from b_i 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 (c, 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 4.12: *In all finite games,*

- (a) *every sequential equilibrium is a Nash equilibrium;*
- (b) *in games of perfect recall, a strategy profile c 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.

4.5 Interim Sequential Equilibrium

As we said, we view our notions of quasi-perfect equilibrium and sequential equilibrium as *ex ante* notions. The players decide on a mixed strategy at the beginning of the game, and do not get to change it. Each player i makes her decision in such a way that it will be optimal conditional on reaching each of her information sets (or conditional on reaching any one of a set of her information sets).

It seems perfectly reasonable to consider interim notions of quasi-perfect equilibrium and sequential equilibrium as well, where the view is that, at each information set for player i , the player reconsiders what to do. We discuss such interim notions here. For the remainder of this discussion, we focus on sequential equilibrium. For simplicity, we also consider only games without absentmindedness, so as to avoid having to deal with questions about how to ascribe beliefs to players at information sets. While there is no controversy about how this should be done in games without absentmindedness, this is not the case in games with absentmindedness (see, for example, [Grove and Halpern 1997]).

Reconsideration at an information set X allows a player i to switch from a strategy b to a strategy b' . If we allow such switches, then we need to be careful to explain whether i remembers that she has switched strategies. If she does not remember that she has switched, then “switching” to a different strategy is meaningless. On the other hand, as we observed in the introduction, allowing the player to remember the switch is in general incompatible with the exogenously-given information structure. For example, if the agent can remember the switch from b to b' in the match-nature game, she is effectively using a “strategy” that makes different moves at x_3 and x_4 .

There are a number of ways of dealing with this problem. The first is to restrict changes at X to strategies of the form $[b, X, b']$. That is, we can simply use Definition 4.5 without change. While this solves the problem, the motivation that we gave earlier for restricting to strategies of the form $[b, X, b']$ no longer applies. While *ex ante*, switching to a strategy that makes a player better off at X and worse off at information sets that do not come after X is not an improvement, once the player is at X , there is no reason for her to care what happens at nodes that do not come after X .

If there is a unique optimal strategy conditional on reaching each information set X (and also *ex ante*) given the agent’s belief assessment and the strategy profile of the other players, then we can give another motivation for considering changes only to strategies of the form $[b, X, b']$. In this case, if switching from b to $[b, X, b']$ was an improvement for player i at information set X , and $X \prec X'$, player i does not have to remember what he decided at X ; he can reconstruct it. But at a node x' in an information set X' such that $X \prec' X'$ but $X \not\prec X'$, player i cannot be sure that he actually went through X , and thus cannot be sure that he actually switched strategies. In this case, it is not unreasonable for the player to use the strategy originally chosen. This leads to using strategies of the form $[b, X, b']$.

Battigalli [1997] considers another variant of interim sequential equilibrium that he

calls *constrained time consistency*.⁹ This is even more restrictive than the notion we have considered here; it further restricts the kinds of changes allowed at an information set X . Given a (behavioral) strategy b , define an ordering \prec_b on information sets by taking $X \prec_b X'$ if $X \prec X'$ and $\pi^b(X') = 0$. Given an action a , let the strategy $\langle b, X, a, b' \rangle$ be the strategy that agrees with b except at X , where the action a is played, and at information sets X' such that $X \prec X'$, where $b'(X')$ is played. Battigalli's motivation for considering strategies of the form $\langle b, X, a, b' \rangle$ is somewhat similar to that given in the second argument: if the player reconsiders at X by playing a , he will remember his initial strategy choice b and play it unless he is at an information set that he could not have reached by playing b . This will serve as a signal that he changed strategies, and he will be able to somehow reconstruct the choice of b' . But if there is not a unique optimal strategy for a player conditional on reaching an information set X , it is not clear how this reconstruction will work. Moreover, it is not clear why the deviation at X should be to a specific action rather than a distribution over actions.

All this discussion is intended to show is that the interplay with the exogenously-given structure and how much of the strategy is recalled makes defining an interim notion of sequential equilibrium delicate.

We conclude this section by considering one approach to defining interim equilibrium that is very much in the spirit of how Piccione and Rubinstein seem to be handling their examples. As we show, by defining the details carefully, we are led to considering a different but related game. Moreover, the *ex ante* sequential equilibrium of the related game acts like the interim sequential equilibrium of the original game.

To be more specific, PR seem to assume that, from time to time, the decision maker may reconsider his move. This decision is *not* a function of the information set; if it were, reconsideration would either happen at every point in the information set (and necessarily happen first at the upper frontier), or would not happen at all. *Ex ante* sequential equilibrium captures this situation. Rather, PR implicitly seem to be assuming that, at each node of the game tree, the agent may decide to reconsider his strategy. Moreover, if he does decide to switch strategies, then he will remember his new strategy. We can model this possibility of reconsideration formally by viewing it as under nature's control. For definiteness, we assume that nature allows reconsideration at each node with some fixed probability ϵ . We can model the process of reconsideration by transforming the original game Γ into a *reconsideration game* $\Gamma^{rec,\epsilon}$. We replace each node x where some player i moves in the original game tree by a node x^n where nature moves. With probability $1 - \epsilon$, nature moves to x , where i moves as usual; with probability ϵ , nature moves to x' , where i gets to reconsider his strategy. The game continues as in Γ from x' , with no further reconsideration moves (since we allow reconsideration to happen only once). The information sets in $\Gamma^{rec,\epsilon}$ are determined by the assumption that the agent can recall his strategy if he changes strategies.

Rather than defining the transformation from Γ to $\Gamma^{rec,\epsilon}$ formally, we show how it

⁹Actually, Battigalli [1997] considers only decision problems, not games, but we can easily translate his notion to the context of games.

works in the case of the absentminded driver in Figure 4. Corresponding to the nodes e_1

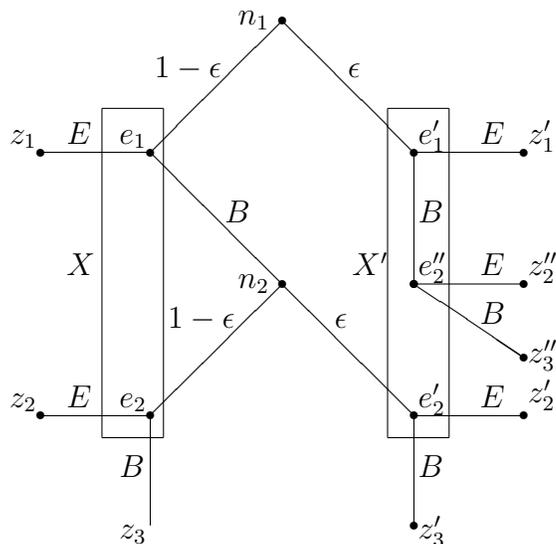


Figure 4: The transformed absentminded-driver game

and e_2 in the original absentminded-driver game, we have moves by nature, n_1 and n_2 . With probability $1 - \epsilon$, we go from n_1 to e_1 , where the driver does not have a chance to reconsider; with probability ϵ , we go to e'_1 . Similarly, n_2 leads to e_2 and e''_2 , with probability $1 - \epsilon$ and ϵ , respectively. From e'_1 , the game continues as before; if he does not exit, the driver reaches the second exit (denoted e''_2), but has no further chance to reconsider. We assume that the driver knows when he has or has had the option of reconsidering, so e'_1 , e'_2 , and e''_2 are in the same information set X' . Implicitly, we are assuming that, because e_1 and e_2 are in the same information set, if the agent decides to do something different at e'_1 , e'_2 , and e''_2 upon reconsideration, he will decide to do the same thing at all these nodes. The nodes e_1 and e_2 from the original game are in information set X . This means that the agent can perform different actions at X' and at X . Call the reconsideration version of the absentminded-driver game Γ^a . Note that the upper frontier of X' consists of e'_1 and e'_2 (although the upper frontier of X consists of just e_1). Moreover, given a strategy b^* , if $\mu_{X'}$ is consistent with b^* , then the probability of e'_i for $i = 1, 2$ is just the normalized probability of reaching e_i under b^* (i.e., $\mu_{X'}(e'_i) = \pi_{b^*}(e_i) / (\pi_{b^*}(e_1) + \pi_{b^*}(e_2))$). As a consequence, a rational agent would use a different action at X and X' , since he would have quite different beliefs regarding the likelihood of being at corresponding nodes in these information sets.

As Piccione and Rubinstein point out, the optimal *ex ante* strategy in the absentminded driver game is to exit with probability $1/3$. But if the driver starts with this strategy and has consistent beliefs, then when he reaches information set X , he will want to exit with probability $2/3$. PR thus argue that there is time inconsistency. In our framework, there is no time inconsistency. As ϵ goes to 0, the optimal *ex ante* strategy in the reconsideration game $\Gamma^{rec,\epsilon}$ (which is also a sequential equilibrium) indeed

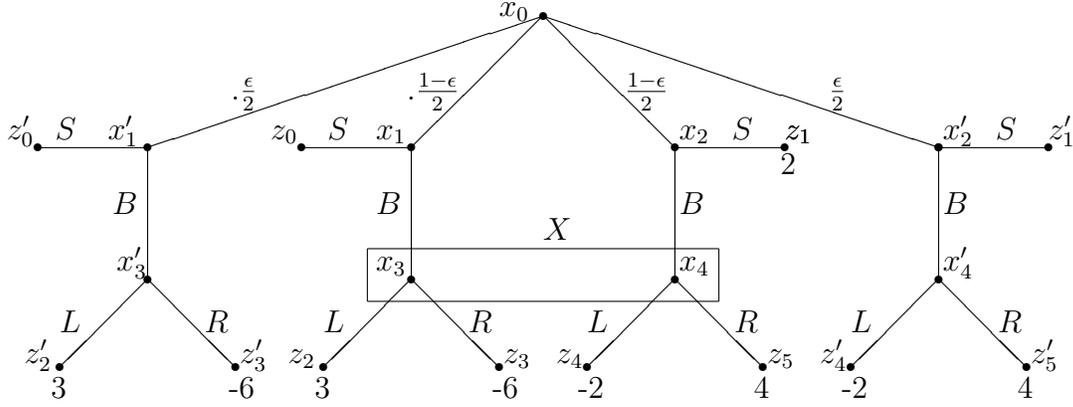


Figure 5: The transformed match-nature game

converges to exiting with probability $1/3$ at nodes in X , and exiting with probability $2/3$ at nodes in X' . But there is nothing inconsistent about this! By capturing the reconsideration process within the game carefully, we can capture interim reasoning, while still maintaining an *ex ante* sequential equilibrium.

We can similarly transform the match-nature game. The result is illustrated in Figure 5, with some slight changes to make it easier to draw. First, we have combined nature's initial "reconsideration" move with the original initial move by nature, so, for example, rather than nature moving to x_1 with probability $\frac{1}{2}$, nature moves to x_1 with probability $\frac{1-\epsilon}{2}$, and to x'_1 , where the agent can reconsider, with probability $\frac{\epsilon}{2}$. For simplicity, we have also omitted the reconsideration at the information set X , since this does not affect the analysis.

Now at the node x'_1 corresponding to x_1 , the agent will certainly want to use the strategy of playing B then L , even though at x_1 he will use the *ex ante* optimal strategy of the original game, and play S (independent of ϵ). Clearly, at both x_2 and x'_2 , he will continue to play B , followed by R . In the reconsideration game, there are four information sets corresponding to the information set X in the original game. There is X itself, the set X' that results from reconsideration at a node in X (which is not shown in the figure), and singleton sets $\{x'_3\}$ and $\{x'_4\}$ that result after reconsideration at x'_1 and x'_2 . We allow x'_3 and x'_4 to be in different information sets because the agent could (and, indeed, will) decide to use different strategies at x'_1 and x'_2 , and hence do different things at x'_3 and x'_4 . Specifically, at x'_1 he will switch to B to be followed by L at x'_3 , while at x'_2 he will continue to use B , to be followed by R at x'_4 . This formalizes the comments that we made in the introduction: the assumption that reconsideration is possible *and that the agent will remember his new strategy after reconsideration* "breaks" the information set $\{x_3, x_4\}$.

Note that every node x in a reconsideration game $\Gamma_{rec,\epsilon}$ can be associated with a unique node in the original game Γ ; we denote this node $o(x)$. In the following discussion, we denote a sequential equilibrium as a pair (c, \cdot) if we want to focus on the strategy profile component. We say that a strategy profile c is a *PR-interim sequential*

equilibrium in a game Γ if, for all ϵ , there exist *ex ante* sequential equilibria (c^ϵ, \cdot) in $\Gamma^{rec, \epsilon}$ such that the strategy profiles c^ϵ converge to c^* , and, for all nodes x in $\Gamma^{rec, \epsilon}$,¹⁰ we have that $c^*(x) = c(o(x))$. The arguments of PR show that there is no PR-interim sequential equilibrium in the absent-minded driver game or the match nature game.

It must be stressed that this approach of using reconsideration games makes numerous assumptions (e.g., an agent remembers his new strategy after switching; nature allows reconsideration at each node with a uniform probability ϵ ; reconsideration happens only once). But, in a precise sense, these assumption do seem to correspond to PR’s arguments.

By way of contrast, in PR’s *modified multiself approach* the agent changes only his action when he reconsiders, and does not remember his new action. We can also model this in our framework using reconsideration games. The structure of the game tree remains the same, but the information sets change. For example, in the reconsideration game corresponding to the absentminded-driver game, the node x''_2 is now in the same information set as x_1 and x_2 ; in the reconsideration game corresponding to the matching nature game, the nodes x'_3 and x'_4 are now in the same information set as x_3 and x_4 . PR show that an *ex ante* optimal strategy is also modified multiself consistent, but in their definition of modified multiself consistent, they consider only information sets reached with positive probability. Marple and Shoham [2013] define a notion of *distributed sequential equilibrium* (DSE) that extends modified multiself consistency to information sets that are reached with probability 0, and prove that a DSE always exists. Taking $\Gamma^{rec', \epsilon}$ to be the reconsideration game appropriate for the modified multiself notion, it is not hard to show that a strategy c is a DSE iff there exist *ex ante* sequential equilibria (b^ϵ, \cdot) in $\Gamma^{rec', \epsilon}$ such the strategies c^ϵ converge to a strategy c^* , and, for all nodes x in $\Gamma^{rec, \epsilon}$, $c^*(x) = c(o(x))$.

This discussion shows that *ex ante* sequential equilibrium can also be a useful tool for understanding interim sequential equilibrium notions.

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.

In this paper, we have defined *ex ante* notions of sequential equilibrium and quasi-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

¹⁰Note that the game tree for $\Gamma^{rec, \epsilon}$ has the same nodes for all choices of ϵ , so it does not matter which ϵ is chosen.

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 example, 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, Halpern, Moses, and Vardi 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.

This is certainly not a new sentiment; work on finite automata playing games, for example, goes back to Neyman [1985] and Rubinstein [1986]. Nevertheless, we believe that there is good reason for describing games by game trees that have perfect recall (but then adding the possibility of imperfect recall later), using an approach suggested by Halpern [1997].

To understand this point, consider a game like bridge. Certainly we may have players in bridge who forget what cards they were dealt, some bidding they have heard, or what cards were played earlier. But we believe that an extensive form description of bridge should describe just the “intrinsic” uncertainty in the game, not the uncertainty due to imperfect recall, where the intrinsic uncertainty is the uncertainty that the player would have even if he had perfect recall. For example, after the cards are dealt, a player has intrinsic uncertainty regarding what cards the other players have. Given the description of the game in terms of intrinsic uncertainty (which will be a game with perfect recall), we can then consider what algorithm the agents use. (In some cases, we may want to consider the algorithm part of the strategic choice of the agents, as Rubinstein [1986] does.) If we think of the algorithm as a Turing machine, the Turing machine determines a *local state* for the agent. Intuitively, the local state describes what the agent is keeping the track of. If the agent remembers his strategy, then the strategy must be encoded in the local state. If he has switched strategies and wants to remember that fact, then this too would have to be encoded in the local state. If we charge the agent for the complexity of the algorithm he uses (as we do in a related paper [Halpern and Pass 2015]), then an agent may deliberately choose not to have perfect recall, since it is too expensive.

The key point here is that, in this framework, an agent can choose to switch strategies, despite not having perfect recall. The strategy (i.e., algorithm) used by the agent determines his information set, and the switch may result in a different information structure. Thus, unlike the standard assumption in game theory (also made in this paper) that information sets are given exogenously, in [Halpern and Pass 2015], the information sets are determined (at least in part) endogenously, by the strategy chosen by the agent. (We can still define exogenous information sets, which can be viewed as giving an upper bound on how much the agent can know, even if he remembers everything.) The

ex ante viewpoint seems reasonable in this setting; before committing to a strategy, an agent considers the best options even off the equilibrium path.¹¹ In [Halpern and Pass 2013], we define sequential equilibrium using the ideas of this paper, adapted to deal with the fact that information sets are now determined endogenously, and show that, again, sequential equilibria exist if we make some reasonable assumptions.

This discussion shows that the appropriate definition of sequential equilibrium in games of imperfect recall—for example, whether we want to use an *ex ante* notion, an interim notion, or something else—will depend in part on the source of imperfect recall. We have argued that an interim-like notion is appropriate for modeling automata playing games, but we need a definition of sequential equilibrium that allows for the information sets being endogenous. Wichardt [2010] presents what can be viewed as another explanation of game trees that exhibit imperfect recall that is quite different in spirit. He assumes that agents do not fully distinguish between different but seemingly similar decisions. Roughly speaking, he models this by grouping nodes where the same moves are available to the agent into one information set. This process can easily result in game trees where the agent can be viewed as having imperfect recall. What would be an appropriate notion of reconsideration and sequential equilibrium in that setting?

As argued by Halpern [1997], modeling an agent’s imperfect recall more carefully, including details like how much of his strategy the agent recalls, may require moving beyond the standard game-tree model that pervades game theory. The use of a different modeling approach may suggest alternative notions of sequential equilibrium. As Marple and Shoham [2013], we can get yet other notions of sequential equilibrium by considering a modified multiself approach, where each player is viewed as being composed of multiple agents (or selves), all sharing the same payoff function, and each controlling a single information set. More research will be needed to understand the space of possible notions of sequential equilibrium and their appropriateness. We believe that the *ex ante* notion that we have defined here will play a key role in that understanding.

Acknowledgments: We thank John Hillas, Jeff Kline, Jörgen Weibull, and anonymous referees of an earlier draft of the paper for very useful comments.

References

- Alpern, S. (1988). Games with repeated decisions. *SIAM Journal of Control and Optimization* 26(2), 468–477.
- Aumann, R. J., S. Hart, and M. Perry (1997). The absent-minded driver. *Games and Economic Behavior* 20, 102–116.
- Battigalli, P. (1997). Dynamic consistency and imperfect recall. *Games and Economic Behavior* 20, 31–50.

¹¹The model does not charge for the *ex ante* consideration. An interim notion of sequential rationality, where we charge for thinking about changes, also would make sense in this setting.

- Fagin, R., J. Y. Halpern, Y. Moses, and M. Y. Vardi (1995). *Reasoning About Knowledge*. Cambridge, MA: 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 J. Y. Halpern (1997). On the expected value of games with absent-mindedness. *Games and Economic Behavior* 20, 51–65.
- Halpern, J. Y. (1997). On ambiguities in the interpretation of game trees. *Games and Economic Behavior* 20, 66–96.
- Halpern, J. Y. and R. Pass (2013). Sequential equilibrium in computational games. In *Proc. Twenty-Third International Joint Conference on Artificial Intelligence (IJCAI '13)*, pp. 171–176.
- Halpern, J. Y. and R. Pass (2015). Algorithmic rationality: Game theory with costly computation. *Journal of Economic Theory* 156, 246–268.
- Hendon, E., H. J. Jacobsen, and B. Sloth (1996). The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior* 12(2), 274–282.
- Hillas, J. and D. Kvasov (2017). Backward induction in games without perfect recall. Unpublished manuscript.
- Isbell, J. (1957). Finitary games. In M. Dresher and H. W. Kuhn (Eds.), *Contributions to the Theory of Games III*, pp. 79–96. Princeton, N.J.: Princeton University Press.
- Kaneko, M. and J. J. Kline (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 R. B. Wilson (1982). Sequential equilibria. *Econometrica* 50, 863–894.
- Kuhn, H. W. (1953). Extensive games and the problem of information. In H. W. Kuhn and A. W. Tucker (Eds.), *Contributions to the Theory of Games II*, pp. 193–216. Princeton, N.J.: Princeton University Press.
- Lipman, B. L. (1997). More absentmindedness. *Games and Economic Behavior* 20(1), 97–101.
- Marple, A. and Y. Shoham (2013). Equilibria in games of imperfect recall. Unpublished manuscript.
- Neyman, A. (1985). Bounded complexity justifies cooperation in finitely repeated prisoner’s dilemma. *Economic Letters* 19, 227–229.
- Okada, A. (1987). Complete inflation and perfect recall in extensive games. *International Journal of Game Theory* 16, 85–91.

- Osborne, M. J. and A. Rubinstein (1994). *A Course in Game Theory*. Cambridge, MA: MIT Press.
- Piccione, M. and A. Rubinstein (1997a). The absent-minded driver's paradox: synthesis and responses. *Games and Economic Behavior* 20(1), 121–130.
- Piccione, M. and A. Rubinstein (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.
- van Damme, E. (1984). A relationship between perfect equilibria in extensive form games and proper equilibria in normal form games. *International Journal of Game Theory* 13, 1–13.
- 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.
- Wichardt, P. C. (2010). Modelling equilibrium play as governed by analogy and limited foresight. *Games and Economic Behavior* 70(2), 472–487.