

Sequential Equilibrium in Games of Imperfect Recall*

JOSEPH Y. HALPERN[†], Cornell University, USA

RAFAEL PASS[‡], Cornell University (Cornell Tech), USA

Although the definition of sequential equilibrium can be applied without change to games of imperfect recall, doing so leads to arguably inappropriate results. We redefine sequential equilibrium so that the definition agrees with the standard definition in games of perfect recall, while still giving reasonable results in games of imperfect recall. The 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. A notion of *interim* sequential equilibrium is also considered.

CCS Concepts: • **Networks** → **Network protocols**; • **Theory of computation** → **Algorithmic game theory and mechanism design**.

Additional Key Words and Phrases: absentmindedness, imperfect recall, sequential equilibrium

ACM Reference Format:

Joseph Y. Halpern and Rafael Pass. 2022. Sequential Equilibrium in Games of Imperfect Recall. *ACM Transactions on Economics and Computation* 1, 1 (February 2022), 26 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

1 INTRODUCTION

Sequential equilibrium [Kreps and Wilson 1982] is one of 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. Solution concepts that take into account what happens off the equilibrium path are important, for example, to eliminate Nash equilibria that arise due to non-credible threats. 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. Consider even a relatively short card game such as bridge. 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

*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, although there are important technical differences between the current version and the preliminary version.

[†]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.

Authors' addresses: Joseph Y. Halpern, Cornell University, Computer Science Department, Ithaca, NY, 14853, USA, halpern@cs.cornell.edu; Rafael Pass, Cornell University (Cornell Tech), Computer Science Department, New York, NY, 10044, USA, rafael@cs.cornell.edu.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Association for Computing Machinery.

Manuscript submitted to ACM

Manuscript submitted to ACM

1

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. Nevertheless, the intuition that underlies these solution concepts, namely, that 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 (at x_1 and x_2), he forgets it if he continues. Thus, this is a game of imperfect recall.




Fig. 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_m 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_m . 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_m) and then switches to b' (playing B at x_1 and L at X_m) 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_m .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

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.

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 *absentminded 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.




Fig. 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 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. ■

Although the definition of sequential equilibrium can be applied without change to games of imperfect recall, doing so leads to arguably inappropriate results. For example, it is not hard to show that there are no sequential equilibria in the match-nature game but, perhaps even more seriously, as we pointed out, the deviations that sequential equilibrium would consider are incompatible with the information structure of the game. In this paper, we propose a definition of sequential equilibrium that coincides with the standard definition in games of perfect recall, while still giving reasonable results in games of imperfect recall. Our definition has some commonalities with one developed contemporaneously

by Hillas and Kvasov [2020a; 2020b]; Section 7 of this paper contains a discussion of both the similarities and the differences. As will be clear, comments from John Hillas based on his work with Kvasov had a significant influence on some of our current definitions.

As is well known, to define sequential equilibrium in games of perfect recall, one difficulty is in determining a player’s beliefs regarding how likely she is to be at each node in an information set off the equilibrium path. This is necessary for dealing with non-credible threats. In games of imperfect recall, determining an agent’s beliefs is difficult even in information sets that lie on the equilibrium path. Indeed, we argue that getting a good definition of sequential equilibrium in games of imperfect recall requires a clear interpretation of the meaning of information sets and the restrictions that they impose on the knowledge and strategies of players.

Consider the issue of defining beliefs at an information set. 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 results in there being no sequential equilibria. (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 them, and that they choose them in such a way that they remain optimal even off the equilibrium path. Our *ex ante* notion of sequential equilibrium does not allow changes that break the information structure of the game. For example, in the match-nature game, we do not allow an agent who starts with strategy b to switch to strategy b' at the node x_2 . We want a strategy that is optimal even off the equilibrium path, but we consider only strategies that respect the information structure of the game.

In games of perfect recall, this *ex ante* notion of sequential equilibrium agrees with what we call here *interim sequential equilibrium*, where agents can potentially change strategies at each information set in a way that may break the information structure (so could, for example, switch from b to b' at x_3 in the match-nature game). Defining a reasonable notion of interim sequential equilibrium in games of imperfect recall raises a number of conceptual issues. We discuss these issues in Section 6, and show how the *ex ante* notion might give some insight into defining an approach that explicitly allows for reconsideration at information sets, and thus can arguably be viewed as capturing interim sequential equilibrium.

There is a final issue that we must confront when dealing 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, neither behavioral strategies nor mixed strategies suffice for Nash equilibrium; we must move to *behavioral strategy mixtures*, that is, distributions over behavioral

strategies [Isbell 1957]. Thus, we work in this paper with behavioral strategy mixtures, which leads to a few additional technical complications.

Both our notions of *ex ante* and interim sequential equilibrium assume that the basic object of choice is a strategy. But as the match-nature game already shows, there are conceptual issues that arise if agents can switch from one strategy to another at an arbitrary information set. Roughly speaking, we deal with this in our *ex ante* notion by restricting the choice to initial information sets, and deal with this in the interim notion by considering a different game related to the initial game, where the switch is in a sense less problematic. Another approach to dealing with this problem is to consider *actions* as the basic object of choice, rather than strategies. This is essentially what is done, for example, in PR's *modified multiselves* approach, and in the approach taken by Aumann, Hart, and Perry [1997]. Lambert, Marple, and Shoham [2019] define a number of equilibrium notions in games of imperfect recall based on this intuition. By taking actions to be the object of choice, they can restrict to behavioral strategies, and do not need to consider behavioral strategy mixtures. While their notions, like ours, agree with the standard definitions in games of perfect recall, they are based on significantly different intuitions (and do not, in general, agree with ours in games of imperfect recall).

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 behavioral strategy mixtures; these preliminaries will be necessary for understanding our formal definition of (*ex ante*) sequential equilibrium. In Section 3, we describe our approach to ascribing beliefs, which also plays a key role in our definitions. In Section 4 we define perfect equilibrium in games of imperfect recall; our definition is identical to Selten's, except that we deal with behavioral strategy mixtures. This forms the basis of our definition of *ex ante* sequential equilibrium, which is given in Section 5. We then discuss interim sequential equilibrium in Section 6. We conclude in Section 7 with a discussion of related work and some related topics.

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 (e.g., see, [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 . The absentminded driver game exhibits absentmindedness; the match-nature game does not.

Note that above, following Osborne and Rubinstein [1994], we spoke of histories in an information set, rather than nodes. We use the words interchangeably, but it worth recalling that, formally, a history is a sequence of actions. Each history leads to a unique node in the game tree. If history h' is a prefix of h (we can talk about prefix here since a history is a sequence), 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.

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, 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 reasonable to assume that the agent commits initially to playing a 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, as we do in Section 6.

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 actions. (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 of imperfect recall without absentmindedness, for every behavioral strategy, there is an outcome-equivalent mixed strategy; however, there is a game of imperfect recall without absentmindedness and a mixed strategy in this game that is not outcome-equivalent to a behavioral strategy [Kuhn 1953]. Once we allow absentmindedness, as pointed out by PR, there may be behavioral strategies that are not outcome-equivalent to any mixed strategy. 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 latter observation also shows that a nontrivial mixture of the two pure strategies is 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 in behavioral strategies. This is no longer the case in games of imperfect recall. Isbell [1957] gives an example of a game with imperfect recall with no Nash equilibrium in behavioral strategies or mixed 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.² 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. We use \mathcal{B}_i^Γ and C_i^Γ to denote the set of player i 's behavioral strategies and behavioral strategy mixtures, respectively, in game Γ .

There are uncountably many behavioral strategies, so a behavioral strategy mixture, which is a probability on behavioral strategies, can be quite a complicated object. It may seem unreasonable to think of resource-bounded players as using them. As observed above, in games of imperfect recall without absentmindedness, we do not need them; it suffices to consider mixed strategies. In games with absentmindedness, however, they seem unavoidable. Fortunately, players have to mix over only finitely 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.³*

A technical consequence of Proposition 2.1 is that in a finite game Γ where player i has d_i information sets and can play at most k_i actions at each one, we can identify a behavioral strategy mixture for player i with an element of $([0, 1] \times [0, 1]^{k_i d_i})^{D_\Gamma}$ —each behavioral strategy mixture 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 behavioral strategy for player i , and thus in $[0, 1]^{k_i d_i}$, for $j = 1, \dots, D_\Gamma$. 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.

²Isbell [1957] actually calls them mixed strategies; Selten [1975] called them *behavior strategy mixtures*; Hillas and Kvasov [2020a; 2020b] call them *general strategies*. We follow Selten, but write “behavioral” rather than “behavior” for consistency with our terminology elsewhere.

³Alpern [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 result and for pointing out that Alpern proved the stronger result.

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 check at each information set whether they are playing optimally by comparing what they are doing to something else that they could have done. Since players can randomize, 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 can 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

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

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 BELIEF SYSTEMS IN GAMES OF IMPERFECT RECALL

Fix a game Γ . Kreps and Wilson [1982] define a *belief system* μ for Γ to be 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 absentminded driver's game, the *ex ante* optimal strategy reaches e_1 with probability 1 and reaches e_2 with probability 2/3.

Definition 3.1. [Halpern 1997] Given an information set X , the *upper frontier of* X , denoted \hat{X} , consists 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 $\hat{X} = X$. This is not true with absentmindedness; for example, $\hat{X}_e = \{e_1\}$.

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 , is 1. Since $\hat{X} = X$ in games of perfect recall, this requirement generalizes that of Kreps and Wilson. Moreover, the requirement holds if we define μ_X in the obvious way:

CLAIM 3.2. *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, $\sum_{x \in \hat{X}} \mu_X(x) = \sum_{x \in \hat{X}} \pi_c(x | X) = \pi_c(\hat{X} | X) = 1$. ■

Having a belief system sufficed for Kreps and Wilson to define the probability on terminal histories given a behavioral strategy, conditional on reaching an information set X . Unfortunately, it does *not* suffice if we are given a behavioral strategy mixture c , since reaching X may cause an agent to update her beliefs regarding which behavioral strategies the other agents chose initially. The following example illustrates this issue.

Example 3.3. Consider a game where player 1 moves either left or right (L or R) at x_0 , then player 2 moves either left or right (lt or rt) at nodes x_1 and x_2 , and then player 1 moves either left or right (ℓ or r) at nodes x_3 – x_6 , which form an information set for player 1. Ignoring the information sets, the game is illustrated in Figure 3; players are assumed to get identical payoffs. First suppose (it is commonly known that) that player 1 uses the behavioral strategy mixture c_1 of playing $L; \ell$ with probability $1/2$ and playing $R; r$ with probability $1/2$, and that player 2 learns player 1’s first move (so that $\{x_1\}$ and $\{x_2\}$ are separate information sets). Clearly at x_1 player 2 should be sure that player 1 initially chose $L; \ell$ while at x_2 player 2 should be sure that player 1 chose $R; r$.

Next suppose that player 1 still plays c_1 , but player 2 does not learn player 1’s first move, so now we have an information set $X = \{x_1, x_2\}$. At X , player 2 should ascribe probability $1/2$ to each of x_1 and x_2 . But player 2 continues to know that if the actual node is x_1 , then player 1 must have chosen $L; \ell$ and if the actual node is x_2 , then player 1 must have chosen $R; r$. As a consequence, given the payoffs, player 2 should play rt .

Finally, suppose that player 1’s strategy is $c_2 = \frac{2}{3}b_1 + \frac{1}{3}b_2$, where b_1 is $.9(L; \ell) + .1(R; r)$ while b_2 is $.1(L; \ell) + .9(R; r)$. Again, player 2 does not learn player 1’s first move. A straightforward computation shows that x_1 is reached with probability $\frac{2}{3} \times .9 + \frac{1}{3} \times .1 = \frac{19}{30}$ and x_2 is reached with probability $\frac{2}{3} \times .1 + \frac{1}{3} \times .9 = \frac{11}{30}$. Moreover, note that x_1 is 18 more times as likely to be reached via b_1 than b_2 , so at x_1 , player 2 should believe that player 1 is playing b_1 with probability $\frac{18}{19}$ and b_2 with probability $\frac{1}{19}$. Similarly, at x_2 , player 2 should believe that player 1 is playing b_1 with probability $\frac{2}{11}$ and b_2 with probability $\frac{9}{11}$. ■

As Example 3.3 shows, agent i ’s beliefs at an information set X involve not only a probability on the nodes in the information set, but an association of each node $x \in X$ with a probability on behavioral strategy profiles. There is an added complication in defining belief systems in our setting. For the *ex ante* notion of sequential equilibrium, we think of calculating these beliefs initially, but after each agent i has chosen which behavioral strategy in her behavioral strategy mixture she will play; moreover, we assume that i will remember her choice of behavioral strategy. This choice clearly affects her beliefs at an information set. These observations underlie the following definition.

Definition 3.4. A *generalized belief system* for a game Γ is a pair (μ, ν) , where μ associates with each information set X for player i and behavioral strategy b_i for i a probability $\mu_{b_i, X}$ on the nodes in X such that $\sum_{x \in \hat{X}} \mu_{b_i, X}(x) = 1$, and ν




Fig. 3. Computing beliefs at information sets.

associates with each node $x \in X$ and behavioral strategy b_i a probability $v_{b_i, x}$ on behavioral strategy profiles in \mathcal{B}_{-i} with finite support.⁴ ■

In the spirit of Kreps and Wilson, we will be particularly interested in generalized belief systems that are consistent with a profile c of behavioral strategy mixtures. In this case, the probability $\mu_{b_i, X}(x)$ is the probability that a node $x \in X$ is reached given that i chose strategy b_i initially (we are interested only if b_i is in the support of c_i) and for b'_{-i} in the support of c_{-i} , $v_{b_i, x}(b'_{-i})$ is the relative probability that x is reached by (b_i, b'_{-i}) (relative to other behavioral strategies in the support of c_{-i}). Clearly, i should not assign positive probability to a strategy profile that would not get chosen initially. As suggested by Example 3.3, it is straightforward to calculate these beliefs if X is reachable by (b_i, c_{-i}) . To deal with information sets that are not reachable by (b_i, c_{-i}) , we proceed in the same spirit as Kreps and Wilson.

Definition 3.5. A behavioral strategy b_i for player i is *completely mixed* if, for each information set X for i and action a that can be played at 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. ■

For each completely mixed strategy c' we can define a generalized belief system $(\mu^{c'}, v^{c'})$ by conditioning in the obvious way (as we did in Example 3.3): if b_i is in the support of c'_i , then $\mu_{b_i, X}(x)$ is just $\pi_{b_i, c'_{-i}}(x | X)$, and if b'_{-i} is in the support of c'_{-i} , then $v_{b_i, x}(b'_{-i})$ is just the relative probability of reaching x with strategy profile (b_i, b'_{-i}) ; more precisely, $v_{b_i, x}(b'_{-i}) = \frac{c'_i(b_{-i})\pi_{(b_i, b'_{-i})}(x|X)}{\pi_{(b_i, c'_{-i})}(x|X)}$.

The idea is to take (μ, v) to be consistent with c if (μ, v) is the limit of (μ^{c^n}, v^{c^n}) for some sequence c^n of completely mixed behavioral strategy mixtures converging to c . The problem is that to define $\mu_{b_i, X}$, we need to consider some sequence $\mu_{b_i^n, X}$, where b_i^n is in the support of c_i^n . Which behavioral strategy b_i^n do we take? If a sequence c_i^1, c_i^2, \dots of behavioral strategy mixtures converges to c_i , then for each behavioral strategy b_i in the support of c_i , there exists a set B_i^n of behavioral strategies in the support of c_i^n such that the sets B_i^n converge to b_i (i.e., for all $\epsilon > 0$, there exists n such that for all $n' > n$, all the strategies in $B_i^{n'}$ are within ϵ of b_i^j , where we can measure the distance between two behavioral strategies b and b' for agent i by taking the sum, over all of i 's information sets X , of the distances

⁴In earlier versions of this paper, we considered a belief assessment that was just like that of Kreps and Wilson; that is, it had only the μ component. We thank John Hillas for comments that emphasized the need to consider agent i 's initial choice of behavioral strategy b_i and the need for i to update her beliefs regarding the probability of the behavioral strategies chosen by other players, as captured by v , similar in spirit to what was done in his joint work with Kvasov [Hillas and Kvasov 2020a,b].

between the probability measures on actions at X determined by b and b' and $c_i^n(B_i^n)$ converges to $c_i(b_i)$. We call B_i^n the *analogue of b_i in c_i^n* . When computing $\mu_{b_i, X}$, we thus consider $\mu_{b_i^n, X}$ for all strategies $b_i^n \in B_i^n$ and weight these beliefs according to the relative weight of b_i^n in B_i^n . We proceed similarly in computing ν .

Definition 3.6. A generalized belief system (μ, ν) 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, if (μ^n, ν^n) is the generalized belief system determined by c^n as above, then for each behavioral strategy b_i in the support of c_i and information set X for i , $\mu_{b_i, X}(x) = \lim_{n \rightarrow \infty} \sum_{b_i^n \in B_i^n} \frac{c_i^n(b_i^n)}{c_i^n(B_i^n)} \mu_{b_i^n, X}^n(x)$, where B_i^n is the analogue of b_i in c_i^n , and if b'_{-i} is a behavioral strategy profile in the support of c_{-i} , then $\nu_{b_i, X}(b'_{-i}) = \lim_{n \rightarrow \infty} \sum_{b_i^n \in B_i^n} \frac{c_i^n(b_i^n)}{c_i^n(B_i^n)} \nu_{b_i^n, X}^n((B'_{-i})^n)$, where $(B'_{-i})^n$ is the analogue of b'_{-i} in c_{-i}^n . ■

Note that if (μ, ν) is consistent with c and X is an information set for i , then we care about $\mu_{b_i, X}$ and $\nu_{b_i, X}$ only if b_i is in the support of c_i . In that case, $\nu_{b_i, X}$ can be viewed as determining a behavioral strategy mixture profile c'_{-i} whose support is a (not necessarily strict) subset of that of c_{-i} . If c is behavioral strategy profile b , then ν is trivial: $\nu_{b_i, X}(b_{-i}) = 1$ for all $x \in X$. Thus, there is no need for ν if we restrict to behavioral strategies (as Kreps and Wilson do).

Given a generalized belief system (μ, ν) and an information set X , like Kreps and Wilson, we want to define a probability distribution $P_{b'_i}^{b_i, \mu, \nu, X}$ on terminal histories that can be thought of as the probability that arises when starting from X with the beliefs on X determined by $\mu_{b_i, X}$ and the beliefs on strategies played by the players other than i at $x \in X$ determined by $\nu_{b_i, X}$, and then continuing from X with player i switching to b'_i (while the players other than i play the behavioral strategy mixture profile determined by $\nu_{b_i, X}$). For each terminal history z , if there is no prefix of z in X , then $P_{b'_i}^{b_i, \mu, \nu, X}(z) = 0$; otherwise, if x_z is the shortest history in X that is a prefix of z , then $P_{b'_i}^{b_i, \mu, \nu, X}(z)$ is the product of $\mu_{b_i, X}(x_z)$ and the probability that (b'_i, c'_{-i}) leads to the terminal history z when started in x_z , where c'_{-i} is the behavioral strategy mixture determined by $\nu_{b_i, X}$.

If c is a behavioral strategy b and (μ, ν) is consistent with c , then our definition of $P_{b'_i}^{b_i, \mu, \nu, X}$ is essentially equivalent to the definition of $P^{\mu, (b'_i, b_{-i})}(\cdot | X)$ given by Kreps and Wilson [1982] for games of perfect recall. As we observed, in that case, $\nu_{b_i, X}(b_{-i}) = 1$ for all $x \in X$, so we can replace ν by b_{-i} . The only difference is that in games of perfect recall, a terminal history has at most one prefix in X . This is no longer the case in games of imperfect recall, so we must specify which prefix to select. Note that if a terminal history z has a prefix in X , then the shortest prefix of z in X is in \hat{X} .

The next result is a sanity check; if we reach X with positive probability, then we can use $P_{b'_i}^{b_i, \mu, \nu, X}$ to calculate the expected probability.

PROPOSITION 3.7. *If (μ, ν) is consistent with c , b_i is in the support of c_i , X is an information set for player i that is reached with positive probability by (b_i, c_{-i}) , and b'_i is a strategy that agrees with b_i at every information set that includes a node that precedes a node in X , then $P_{b'_i}^{b_i, \mu, \nu, X} = \pi_{b'_i, c_{-i}} | X$.*

Proof: Suppose that z is a terminal history. Clearly if z has no prefix in X , then $P_{b'_i}^{b_i, \mu, \nu, X}(z) = \pi_{(b'_i, c_{-i})}(z | X) = 0$. If z has a prefix in X , let x_z denote the unique prefix in \hat{X} . Given an arbitrary behavioral strategy profile b'' , let $\pi_{b'', x_z}(z)$ denote the probability of reaching z when starting at x_z and playing b'' . Note that $\pi_{b''}(z) = \pi_{b''}(x_z) \pi_{b'', x_z}(z)$. Suppose that c_{-i} has support $b_{-i}^1, \dots, b_{-i}^k$ and that $c_{-i}(b_{-i}^j) = \alpha_j$. Then, since $\mu_{b_i, X}(x_z) = \pi_{b_i, c_{-i}}(x_z | X) = \pi_{b'_i, c_{-i}}(x_z | X)$, we

have

$$\begin{aligned}
\pi_{(b'_i, c_{-i})}(z | X) &= \alpha_1 \pi_{(b'_i, b_{-i}^1)}(x_z | X) \pi_{(b'_i, b_{-i}^1), x_z}(z) + \dots + \alpha_k \pi_{(b'_i, b_{-i}^k)}(x_z | X) \pi_{(b'_i, b_{-i}^k), x_z}(z) \\
&= \pi_{b_i, c_{-i}}(x_z | X) \left(\frac{\alpha_1 \pi_{(b_i, b_{-i}^1)}(x_z | X)}{\pi_{b_i, c_{-i}}(x_z | X)} \pi_{(b'_i, b_{-i}^1), x_z}(z) + \dots + \frac{\alpha_k \pi_{(b_i, b_{-i}^k)}(x_z | X)}{\pi_{b_i, c_{-i}}(x_z | X)} \pi_{(b'_i, b_{-i}^k), x_z}(z) \right) \\
&= \mu_{b_i, X}(x_z) (v_{b_i, x_z}(b_{-i}^1) \pi_{(b'_i, b_{-i}^1), x_z}(z) + \dots + v_{b_i, x_z}(b_{-i}^k) \pi_{(b'_i, b_{-i}^k), x_z}(z)) \\
&= P_{b'_i}^{b_i, \mu, v, X}(z),
\end{aligned}$$

as desired. ■

4 PERFECT EQUILIBRIUM

We start by considering perfect equilibrium in games of imperfect recall, since our definition is quite straightforward. Indeed, we use literally the same definition as Selten [1975], except that we use behavioral strategy mixtures rather than behavioral strategies. The first step is to slightly extend Selten's notion of *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 $\eta > \vec{0}$ (i.e., $\eta_a > 0$ for all actions a), then $\pi_c(X) > 0$ for all information sets X , and if $\eta = \vec{0}$ then the acceptable strategy profiles in (Γ, η) are just the strategy profiles in Γ .

Definition 4.2. The behavioral strategy mixture profile c^* is a *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) $c^k \rightarrow c^*$.⁵ ■

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 4.3. *A 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

⁵We 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.

the fact that, by Proposition 4.1, the set of behavioral strategy mixtures acceptable in (Γ, η_k) is compact. Moreover, as observed after the proof of Proposition 2.1, the set of behavioral strategy mixtures in a finite game Γ can be identified with the compact set $([0, 1] \times [0, 1]^{k_i d_i})^{D_i}$. Thus, 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 perfect equilibrium. ■

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.

As we noted above, a 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 (see Proposition 5.3).

5 DEFINING SEQUENTIAL EQUILIBRIUM

In this section, we define the notion of sequential equilibrium in games of imperfect recall. Again, we try to stick closely to the spirit of the original definition. To do that, we first carefully define (in Section 5.1) what it means to switch from a behavior strategy b to b' at an information set X ; we must do so in a way that ensures that we continue to have a well-defined strategy (where the same action is performed at all nodes in an information set). Once we make this precise, we will be able to define a notion that we call *sequential' equilibrium* (see Section 5.2). It turns out that strategies b and b' in the match-nature game are both sequential' equilibria. Thus, sequential' equilibrium would not be a refinement of Nash equilibrium. We will deal with this problem by considering changes not just at one information set, but at several; see Section 5.3.

5.1 The strategy $[b, X, b']$

If b and b' are behavioral strategies for player i , we would like to define $[b, X, b']$ as 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 .

We now define a notion of “after” for information sets, and use that to define $[b, X, b']$ in a way that guarantees that it is a strategy.

Definition 5.1. For nodes x and x' , write $x < x'$ if x precedes x' in the game tree (i.e., if the history leading to x is a prefix of that leading to x'). Extend $<$ to a partial order on information sets by defining $X < X'$ iff, for all $x' \in X'$, there exists some $x \in X$ such that $x < x'$. Define $x \leq x'$ iff $x = x'$ or $x < x'$; similarly, $X \leq X'$ iff $X = X'$ or $X < X'$. Finally, if b and b' are behavioral strategies for player i , and X is an information set for i , then $[b, X, b']$ is the behavioral strategy according to which i plays b' at every information set X' for i such that $X \leq X'$, and otherwise plays b .⁶ ■

⁶We do not define $[c, X, c']$ if either c or c' is a behavioral strategy mixture (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_m , since we do not have $\{x_1\} \leq X_m$. Thus, $[b, \{x_1\}, b']$ as we have defined it is not better for the player than 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 $x' \leq x$ 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 5.2. *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: Clearly if $(b, X, b') = [b, X, b']$, then (b, X, b') is a behavioral strategy (since $[b, X, b']$ is). 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 $X \leq X'$, then at every node $x \in X'$, the player plays $b'(X')$ at x according to both (b, X, b') and $[b, X, b']$. Thus, it must be the case that $X \not\leq X'$. 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 $X \not\leq X'$, 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. ■

Before going on with sequential equilibrium, we can now make precise the sense in which a perfect equilibrium is a best response at each information set. The following notation will prove useful here and in the sequel. If (μ, ν) is consistent with the behavioral strategy mixture profile c , b_i is in the support of c_i , X is an information set for player i , and b'_i is a behavioral strategy for player i , let $EU_i([b_i, X, b'_i], c_{-i} \mid b_i, \mu, \nu, X)$ denote the expected utility for player i with respect to $P_{[b'_i, X, b'_i]}^{b_i, \mu, \nu, X}$; that is,

$$EU_i([b_i, X, b'_i], c_{-i} \mid b_i, \mu, \nu, X) = \sum_{z \in Z} P_{[b'_i, X, b'_i]}^{b_i, \mu, \nu, X}(z) u_i(z).$$

Intuitively, $EU_i([b_i, X, b'_i], c_{-i} \mid b_i, \mu, \nu, X)$ is player i 's expected utility if he switches from b_i to b'_i at X , while the other players are playing c_{-i} .

PROPOSITION 5.3. *If the behavioral strategy mixture profile c is a perfect equilibrium in game Γ , then there exists a generalized belief system (μ, ν) such that for all player i , all behavioral strategies b_i in the support of c_i , all information sets X for player i , and all behavioral strategies b'_i for player i , we have*

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

Proof: Since c is a 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) . All the behavioral strategies in the support of c^n are completely mixed (since they are strategy profiles in perturbed games).

Let (μ^n, ν^n) be the generalized belief assessment consistent with c^n , as in Definition 3.6. We can assume without loss of generality that $\lim_{n \rightarrow \infty} (\mu^n, \nu^n)$ exists. (Since Γ is a finite game, by Proposition 2.1, we can find a subsequence of c^1, c^2, \dots for which the limit exists, and we can replace the original sequence by the subsequence.) Let (μ, ν) be this limit.

We claim that the result holds with respect to (μ, ν) . For suppose not. Then there exists a player i , an information set X for i , a behavioral strategy b_i in the support of c_i , a behavioral strategy b'_i for i , and $\epsilon > 0$ such that $EU_i((b_i, c_{-i}) \mid b_i, \mu, \nu, X) + \epsilon < EU_i((b_i, X, b'_i), c_{-i}) \mid b_i, \mu, \nu, X)$. Let B_i^n be the analogue of b_i in c^n . Then for all $b_i^n \in B_i^n$, it follows from Proposition 3.7 that $EU_i((b_i^n, c_{-i}^n) \mid X) = EU_i((b_i^n, c_{-i}^n) \mid b_i^n, \mu, \nu, X)$ and that $EU_i((b_i^n, X, b'_i), c_{-i}^n) \mid X) = EU_i((b_i^n, X, b'_i), c_{-i}^n) \mid b_i^n, \mu, \nu, X)$. Thus, $EU_i((b_i^n, c_{-i}^n) \mid X) \rightarrow EU_i((b_i, c_{-i}) \mid b_i, \mu, \nu, X)$ and $EU_i((b_i^n, X, b'_i), c_{-i}^n) \mid X) \rightarrow EU_i((b_i, X, b'_i), c_{-i}) \mid b_i, \mu, \nu, X)$. Since $(b_i^n, c_{-i}^n) \rightarrow (b_i, c_{-i})$ and $\eta_n \rightarrow \vec{0}$, there exists some strategy b'' and $k > 0$ such that b'' is acceptable for $(\Gamma, \eta_{n'})$ for all $n' > n$, and $EU_i((b_i^{n'}, c_{-i}^{n'}) \mid X) + \epsilon/2 < EU_i((b_i^{n'}, X, b''), c_{-i}^{n'}) \mid X)$. But this contradicts the assumption that $c^{n'}$ is a Nash equilibrium of $(\Gamma, \eta_{n'})$ (since this assumption also implies that $EU_i(c^{n'}) = EU_i((b_i^{n'}, c_{-i}^{n'}))$).⁷ ■

We are implicitly identifying “ b_i is a best response for i at information set X ” with “ $EU_i((b_i, c_{-i}) \mid b_i, \mu, \nu, X) \geq EU_i((b_i, X, b'), c_{-i}) \mid b_i, \mu, \nu, X)$ ” 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 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 5.3.

5.2 Sequential' Equilibrium

Recall that Krep and Wilson [1982] defined a sequential equilibrium to be a pair (b, μ) , where b is a behavior strategy, μ is a belief assessment consistent with b , and b is *sequentially rational*, in the sense that for all players i and all information sets X for players i , b_i is a best response to b_{-i} .

We now give a preliminary definition of sequential equilibrium. Our definition is identical to that of Krep and Wilson [1982] for profiles of behavioral strategies. The complexity comes in dealing with profiles of behavioral strategy mixtures. We call this notion *sequential' equilibrium*, and view it as preliminary, because it does not have a critical property: as we show by example, it does not always imply Nash equilibrium in games of imperfect recall. We generalize the definition to get what we believe is a more appropriate notion of sequential equilibrium in the next section. But that definition is more complicated than the one below, and the one below captures most of the essential intuitions, which is why we consider it first.

Definition 5.4. A pair $(c, (\mu, \nu))$ consisting of a behavioral strategy mixture c and a generalized belief assessment (μ, ν) is a *sequential' equilibrium* if (μ, ν) is consistent with c and for all players i , all information sets X for player i , all

⁷Note 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 would happen twice during the game).

behavioral strategies b_i in the support of c_i , and all behavioral strategies b'_i for player i , we have that

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

■

As we have observed, if c is a behavioral strategy, then we do not need the ν component, and do not need to condition on b_i . Thus, our definition can be viewed as a generalization of that of Kreps and Wilson to the setting of imperfect recall, where we must deal with behavioral strategy mixtures.

It is immediate from Proposition 5.3 that every perfect equilibrium is a sequential' equilibrium. Thus, a sequential' equilibrium exists for every game.

COROLLARY 5.5. *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 generalized belief assessment (μ, ν) consistent with the *ex ante* optimal strategy assigns probability 1 to reaching e_1 . For all strategies b , $EU(b \mid b, \mu, \nu, X_e) = EU(b \mid e_1) = EU(b)$, 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, b' is easily seen to be a sequential' equilibrium, since $[b', \{x_2\}, b''] = [b', \{x_1\}, b''] = b'$ for all strategies b'' ; no nontrivial deviation is possible at x_1 and x_2 . However, b' is clearly not a Nash equilibrium; b has a higher expected utility. 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 $<$ 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 5.6. *In all finite games,*

- (a) *every 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 et al. 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 Theorems 5.5 and 5.6(c), this is also the only sequential' equilibrium of the modified game. However, replacing b by b' gives a strategy that satisfies the one-step deviation principle but is not a sequential' equilibrium of the modified game.

5.3 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 $\{x_1, x_2\}$. 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 generalized belief system. Recall that a generalized belief system (μ, ν) for a game Γ associates with each information set X for player i and behavioral strategy for player i in Γ a probability $\mu_{b_i, X}$ on the histories in X . We take an *extended belief system* for a game Γ to be a pair (μ, ν) . Just as in a generalized belief system, for a node x where i moves and behavioral strategy b_i , we take $\nu_{b_i, x}$ to be a probability on behavioral strategies in \mathcal{B}_{-i} with finite support. The μ component of an extended belief system associates with each (non-empty) set \mathcal{X} of information sets for player i and behavioral strategy b_i a probability $\mu_{b_i, \mathcal{X}}$ on the histories in the union of the information sets $X_i \in \mathcal{X}$ such that $\sum_{x \in \tilde{\mathcal{X}}} \mu_{b_i, \mathcal{X}}(x) = 1$, where $\tilde{\mathcal{X}}$ denotes the upper frontier of \mathcal{X} , that is, all the nodes $x \in \cup \mathcal{X}$ such that there is no node $x' \in \cup \mathcal{X}$ with $x' < x$. As before, we interpret $\mu_{b_i, \mathcal{X}}(x)$ as the probability of reaching x conditional on reaching \mathcal{X} .

We define what it means for an extended belief system (μ, ν) to be consistent with a behavioral strategy mixture profile c just as we defined it in the case of a generalized belief system, except that we replace the X in the definition of $\mu_{b_i, X}$ by \mathcal{X} ; the definition of $\nu_{b_i, x}$ remains unchanged. We can also define the probability $P_{b_i}^{b_i, \mu, \nu, \mathcal{X}}$ and $EU_i((b_i, X, b'), c_{-i} | b_i, \mu, \nu, \mathcal{X})$ analogously to before.

The analogue of Proposition 3.7 holds.

PROPOSITION 5.7. *If (μ, ν) is consistent with c , b_i is in the support of c_i , \mathcal{X} is a set of information sets for player i such that $\tilde{\mathcal{X}}$ is reached with positive probability by (b_i, c_{-i}) , and b'_i is a strategy that agrees with b_i at every information set that includes a node that precedes a node in $\cup \mathcal{X}$, then $P_{b'_i}^{b_i, \mu, \nu, \mathcal{X}} = \pi_{b'_i, c_{-i}} | \mathcal{X}$.*

Proof: The proof is identical to the proof of Proposition 3.7. ■

We say that \mathcal{X} *precedes* \mathcal{X}' , written $\mathcal{X} \leq \mathcal{X}'$, iff for all $x' \in \mathcal{X}'$ there exists some $x \in \cup \mathcal{X}$ such that x precedes x' on the game tree; that is, we are defining \leq 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 $\mathcal{X} \leq X'$, and otherwise plays b . We now have the following generalization of Proposition 5.3.

PROPOSITION 5.8. *If the behavioral strategy mixture profile c is a perfect equilibrium in game Γ , then there exists an extended belief system (μ, ν) such that for all player i , all behavioral strategies b_i in the support of c_i , all sets \mathcal{X}_i of information sets for player i , and all behavioral strategies b'_i for player i , we have*

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

Proof: The proof is identical to that of Proposition 5.3, except that we use Proposition 5.7 instead of Proposition 3.7. ■

We now define sequential equilibrium just as we did sequential' equilibrium, but we use sets of information sets rather than individual information sets.

Definition 5.9. A pair $(c, (\mu, \nu))$ consisting of a behavioral strategy mixture c and an extended belief assessment (μ, ν) is a *sequential equilibrium* if (μ, ν) is consistent with c and for all players i , all sets \mathcal{X} of information sets for player i for player i , all behavioral strategies b_i in the support of c_i , and all behavioral strategies b'_i for player i , we have that

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

■

It is immediate from Proposition 5.8 that every perfect equilibrium is a sequential equilibrium. Thus, every game has a sequential equilibrium.

THEOREM 5.10. *A sequential equilibrium exists in all finite games.*

Clearly, every strategy that is part of a sequential equilibrium is part of a sequential' equilibrium. Furthermore, as the definition of sequential equilibrium considers changes at all sets of information sets for all player i , and in particular, the set consisting of all information sets for player i , it follows that every strategy that is part of a sequential equilibrium is a Nash equilibrium. (Recall that this was not the case for sequential' equilibrium.) Finally, suppose that $(c, (\mu, \nu))$ is a sequential' equilibrium of a game Γ of perfect recall. Then we claim that we can extend (μ, ν) to an extended belief system (μ', ν) such that $(c, (\mu', \nu))$ is a sequential equilibrium: Consider the sequence of strategy profiles c^1, c^2, \dots converging to c that determine (μ, ν) ; this sequence also determines an extended belief system (μ', ν) . We claim that $(c, (\mu', \nu))$ 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 extended belief assessment (μ', ν) . 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', ν) . 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 generalized belief assessment (μ, ν) , so $(c, (\mu, \nu))$ cannot be a sequential' equilibrium. We conclude that in games of perfect recall, every behavioral strategy mixture profile that is part of a sequential' equilibrium is also part of 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 part of a sequential' equilibrium, but is not part of 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 5.11. *In all finite games,*

- (a) if $(c, (\mu, \nu))$ is a sequential equilibrium then c is a Nash equilibrium;
- (b) in games of perfect recall, a strategy profile c is part of a sequential equilibrium iff c is part of a sequential' equilibrium.

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

6 INTERIM SEQUENTIAL EQUILIBRIUM

As we said, we view our notion of sequential equilibrium as an *ex ante* notions. Each player i decides on her strategy at the beginning of the game and does not get to change it. 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), given the choices of the other players. The strategy chosen must be consistent with the information structure of the game.

It seems perfectly reasonable to consider interim notions of sequential equilibrium, where the view is that, at each of her information sets, player i reconsiders what to do. The result of this reconsideration may not respect the information structure of the game. After all, why should a player at node x_2 of the match-nature game care about respecting the structure of the game? At x_2 , he knows that he will not reach x_3 .

In this section, we consider one approach to defining interim equilibrium that is very much in the spirit of how Piccione and Rubinstein handle their examples. Before going on, we should note that if we allow reconsideration in games of imperfect recall, we must assume that i remembers that she has switched strategies. If she does not remember that she has switched, then “switching” to a different strategy is meaningless. For example, there is no point in switching from b to b' at x_2 in the match-nature game if the player does not remember at x_4 that she has switched. On the other hand, as we observed in the introduction, allowing the player to remember the switch is not consistent with the information structure. In our view, allowing reconsideration in a game of imperfect recall amounts 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 some node in 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 in the game tree 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, for simplicity, we allow reconsideration to happen only once). The nodes in the game tree after a reconsideration move are in a different information set from the corresponding nodes if there was no reconsideration move; this is meant to capture 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 works in the case of the absentminded driver in Figure 4. Corresponding to the nodes e_1 and e_2 in the original absentminded-driver game,




Fig. 4. The transformed absentminded-driver game.

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'_e . 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_e . This means that the agent can perform different actions at X'_e and at X_e . Call the reconsideration version of the absentminded-driver game Γ^a . Note that the upper frontier of X'_e consists of e'_1 and e'_2 (although the upper frontier of X_e consists of just e_1). Moreover, given a behavioral strategy b^* , if μ is consistent with b^* , then $\mu_{b^*, X'_e}(e'_i)$ for $i = 1, 2$ is just the normalized probability of reaching e_i under b^* (i.e., $\mu_{b^*, X'_e}(e'_i) = \pi_{b^*}(e_i) / (\pi_{b^*}(e_1) + \pi_{b^*}(e_2))$). As a consequence, a rational agent might play different actions at X_e and X'_e , since he would have quite different beliefs regarding the likelihood of being at corresponding nodes in these information sets.

As PR 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_e , 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 converges to exiting with probability $1/3$ at nodes in X_e , and exiting with probability $2/3$ at nodes in X'_e . But there is nothing inconsistent about this! If $\epsilon > 0$, then the probability of exiting at X'_e will be greater than $1/3$ and the probability of exiting at X_e will be less than 1 : unless the probability of exiting in X_e is 1, the probability of reaching e'_2 is positive and consequently (by the same reasoning as in Example 1.2), the probability of exiting at X'_e will




Fig. 5. The transformed match-nature game.

be greater than $1/3$. Thus, the probability of exiting at X_e will be less than $1/3$ (to counter the fact that, in the case of reconsideration, we exit more often). 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_m , 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_m in the original game. There is X_m itself, the set X_m' that results from reconsideration at a node in X_m (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 $\Gamma^{rec, \epsilon}$,⁸ we have that $c^*(x) = c(o(x))$. The arguments of PR show that there is no PR-interim sequential equilibrium in the absentminded 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 assumptions do seem to correspond to PR's arguments.

⁸Note 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.

By way of contrast, in PR’s *modified multiseif 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 multiseif consistent, but in their definition of modified multiseif consistent, they consider only information sets reached with positive probability. Lambert, Marple, and Shoham [2019] define a notion of *distributed sequential equilibrium* (DSE) that extends modified multiseif 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 multiseif 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 that 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.

7 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 given definitions of *ex ante* notions of sequential equilibrium and perfect equilibrium that agree with the standard definitions in game of perfect recall, but seem more appropriate in games of imperfect recall. We have also pointed out the subtleties in doing so.

As we said in the introduction, there has recently been other work giving definitions of equilibrium notions in games of imperfect recall. Hillas and Kvasos [2020a; 2020b] (HK from now on) give results quite similar to ours. Among other things, they define sequential equilibrium and perfect equilibrium, and observe that sequential equilibrium does not satisfy the one-step deviation principle in games of imperfect recall.

We now briefly compare our notion of sequential equilibrium with that of HK.

- HK have an analogue of v in their notion of a belief system.⁹ In games without absentmindedness, the HK notion of v puts a probability on pure strategy profiles; in games with absentmindedness, it puts a probability on behavioral strategy profiles. In both cases, the HK version of v is indexed by information sets, not nodes. Thus, for two different nodes x and x' in an information set, HK cannot distinguish the different beliefs that an agent would have about another agent’s strategies if the true state is x or x' . As we saw in our earlier examples, this can be useful information.
- In games with absentmindedness, HK also have an analogue of our $\mu_{b,\mathcal{X}}$, except that their analogue is defined only for behavioral strategies b that reach \mathcal{X} with positive probability, while our notion is defined for all behavioral strategies b .
- While we require that the agent’s actual strategy be a best response at each set \mathcal{X} of information sets (or, more precisely the agent’s strategy from \mathcal{X} on), HK require that the strategies that the agent believes he is

⁹Indeed, as we pointed out earlier, it was John Hillas who convinced us that we needed an analogue to their v .

following at each set \mathcal{X} of information sets be a best response at \mathcal{X} . The two approaches are equivalent if \mathcal{X} is reached with positive probability, but may not be equivalent for sets not reached by the strategy in games with absentmindedness, as the following example shows.

- When we consider changing strategies at a set \mathcal{X} of information sets (to see if the current strategy is a best response), we allow changes at all information sets Y that come after some information set $X \in \mathcal{X}$. This requires us to carefully define what it means for one information set to come after another. HK avoid that by allowing changes only at the information sets in \mathcal{X} .
- In games with absentmindedness, whereas we define a completely mixed behavioral strategy mixture to be a (finite) distribution whose support consists of completely mixed behavioral strategies (which are defined the standard way), HK take it to be a distribution whose support consists of *all* behavioral strategies.

Despite these differences, in games without absentmindedness, our definitions of sequential equilibrium are equivalent, although ours is perhaps closer in spirit to the original. However, in games with absentmindedness, the differences in the definitions have some bite.

Example 7.1. Consider the one-person game with absentmindedness shown in Figure 6, where there are two information sets: $\{x_0, x_2\}$ and $\{x_1\}$. The strategy $L; \ell$ is not part of a sequential equilibrium of this game according to




Fig. 6. A game where the HK definition of sequential equilibrium differs from ours.

our definition (since the agent can do better at x_1 by switching from ℓ to r); the only sequential equilibrium according to our definition is $L; r$, with the obvious system of beliefs. However, $L; \ell$ is part of a sequential equilibrium according to the HK definition (given in [Hillas and Kvasov 2020b]), as we understand it, while $L; r$ is not.

The only system of beliefs consistent with $L; \ell$ places probability 1 on $R; \ell$ at x_1 (since the only way to reach x_1 is to play R , and a completely mixed strategy consistent with $L; \ell$ must place probability $1 - \epsilon$ on ℓ and probability ϵ on r .) HK thus require $R; \ell$ to be a best response at x_1 , which indeed it is. By way of contrast, our definition still requires $L; \ell$ to be a best response at x_1 , which it is not. A similar argument shows that $L; r$ is not part of a sequential equilibrium for HK, while it is for us.

Conceptually, we are assuming that an agent recalls his behavioral strategy b (or, more precisely, the behavioral strategy b that was chosen initially from the behavioral strategy mixture), and requires b to be a best response even at sets \mathcal{X} of information sets that cannot be reached by b . If \mathcal{X} is not reachable by b , then the fact that the agent is at \mathcal{X}

treated as a mistake; the agent still assumes that he continues with b . HK, on the other hand, assume that the agent will play a “best approximation” to b from X on. Thus, we assume in this example that the agent will continue to play L ; ℓ at the information set, despite the fact that he must have played R to reach x_1 , while HK assume that he will continue to play R when he reaches the information set again. So these definitions are capturing somewhat different intuitions. ■

As we mentioned earlier, Lambert, Marple, and Shoham [2019] consider a multiself approach, where a different “self” of the agent moves at each information set. Moreover, these selves cannot coordinate. As pointed out by HK, this leads to a philosophically quite different notion of equilibrium (and, as a technical matter, their equilibria are, in general, quite different from ours and those of HK).

While we could continue to investigate various solution concepts in games of imperfect recall to sort out when and whether each is appropriate, we believe that a different approach may be more productive. Our definitions were given in the standard game-theoretic model of extensive-form games with information sets (as are those of HK and Lambert, Marple, and Shoham). 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 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 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). We believe that a deeper understanding of imperfect recall can be gained by considering the runs-and-systems framework. Specifically, the idea is that we start by using game trees that have perfect recall, but then adding the possibility of imperfect recall later, using an approach suggested by Halpern [1997].

To understand how this would work, consider a game like bridge. Certainly we may have players in bridge who forget what cards they were dealt, some of the bidding that 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 choice of algorithm to be 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 2019], 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.¹⁰ In [Halpern and Pass 2019], we define a notion of sequential equilibrium for finite automata playing games using the ideas of this paper, adapted to deal with the fact that information sets are now determined endogenously, and show that, again, *ex ante* sequential equilibria exist if we make some reasonable assumptions, and (under further assumptions) coincide with interim sequential equilibria.

This discussion also suggests 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, which can also be captured in the runs-and-systems framework by taking the appropriate notion of local state, can easily result in game trees where the agent can be viewed as having imperfect recall.

We believe that investigating solution concepts in a more general setting like the runs-and-systems framework, which provides more flexibility in capturing agents’ information, will provide further insight and understanding. We have chosen to use the more standard setting here, where information sets are given exogenously, to be able to relate our definitions to the ones in the literature. In any case, as the discussion in [Halpern and Pass 2019] shows, the key ideas underlying the definitions given here extend naturally to the runs-and-systems framework.

Acknowledgments: We thank Jeff Kline, Jörgen Weibull, and anonymous referees of an earlier draft of the paper for very useful comments. We are particularly grateful to John Hillas, whose comments helped us correct errors in an earlier version of the paper.

REFERENCES

- S. Alpern. 1988. Games with repeated decisions. *SIAM Journal of Control and Optimization* 26, 2 (1988), 468–477.
- R. J. Aumann, S. Hart, and M. Perry. 1997. The absent-minded driver. *Games and Economic Behavior* 20 (1997), 102–116.
- R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. 1995. *Reasoning About Knowledge*. MIT Press, Cambridge, MA. A slightly revised paperback version was published in 2003.
- I. Gilboa. 1997. A Comment on the Absent-Minded Driver Paradox. *Games and Economic Behavior* 20, 1 (1997), 25–30.
- A. J. Grove and J. Y. Halpern. 1997. On the expected value of games with absentmindedness. *Games and Economic Behavior* 20 (1997), 51–65.
- J. Y. Halpern. 1997. On ambiguities in the interpretation of game trees. *Games and Economic Behavior* 20 (1997), 66–96.
- J. Y. Halpern and R. Pass. 2015. Algorithmic rationality: Game theory with costly computation. *Journal of Economic Theory* 156 (2015), 246–268. <https://doi.org/10.1016/j.jet.2014.04.007>
- J. Y. Halpern and R. Pass. 2019. Sequential equilibrium in computational games. *ACM Trans. on Economics and Computation* 7, 2 (2019).
- E. Hendon, H. J. Jacobsen, and B. Sloth. 1996. The one-shot-deviation principle for sequential rationality. *Games and Economic Behavior* 12, 2 (1996), 274–282.
- J. Hillas and D. Kvasov. 2020a. Backward induction in games without perfect recall. *Games and Economic Behavior* 124 (2020), 207–218.
- J. Hillas and D. Kvasov. 2020b. Backward induction in nonlinear games. (2020). Unpublished manuscript.
- J. Isbell. 1957. Finitary games. In *Contributions to the Theory of Games III*, M. Dresher and H. W. Kuhn (Eds.). Princeton University Press, Princeton, N.J., 79–96.
- M. Kaneko and J. J. Kline. 1995. Behavior strategies, mixed strategies, and perfect recall. *International Journal of Game Theory* 24 (1995), 127–145.
- J. J. Kline. 2005. Imperfect recall and the relationships between solution concepts in extensive games. *Economic Theory* 25 (2005), 703–710.
- D. M. Kreps and R. B. Wilson. 1982. Sequential equilibria. *Econometrica* 50 (1982), 863–894.
- H. W. Kuhn. 1953. Extensive games and the problem of information. In *Contributions to the Theory of Games II*, H. W. Kuhn and A. W. Tucker (Eds.). Princeton University Press, Princeton, N.J., 193–216.

¹⁰The model does not charge for the *ex ante* consideration. An interim notion of sequential rationality where we charge for thinking about changes would also make sense in this setting.

- N. S. Lambert, A. Marple, and Y. Shoham. 2019. On equilibria in games of imperfect recall. *Games and Economic Behavior* 113 (2019), 164–185.
- B. L. Lipman. 1997. More absentmindedness. *Games and Economic Behavior* 20, 1 (1997), 97–101.
- M. J. Osborne and A. Rubinstein. 1994. *A Course in Game Theory*. MIT Press, Cambridge, MA.
- M. Piccione and A. Rubinstein. 1997a. The absent-minded driver’s paradox: synthesis and responses. *Games and Economic Behavior* 20, 1 (1997), 121–130.
- M. Piccione and A. Rubinstein. 1997b. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20, 1 (1997), 3–24.
- R. T. Rockafellar. 1970. *Convex Analysis*. Princeton University Press, Princeton, N.J.
- A. Rubinstein. 1986. Finite automata play the repeated prisoner’s dilemma. *Journal of Economic Theory* 39 (1986), 83–96.
- R. Selten. 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *International Journal of Game Theory* 4 (1975), 25–55.
- P. C. Wichardt. 2010. Modelling equilibrium play as governed By analogy and limited foresight. *Games and Economic Behavior* 70, 2 (2010), 472–487.