



On equilibria in games with imperfect recall [☆]

Nicolas S. Lambert ^{a,*}, Adrian Marple ^b, Yoav Shoham ^c



^a Graduate School of Business, Stanford University, United States of America

^b Google, Inc., United States of America

^c Department of Computer Science, Stanford University, United States of America

ARTICLE INFO

Article history:

Received 26 October 2013

Available online 24 September 2018

JEL classification:

C70

C72

Keywords:

Imperfect recall

Multiselves

Absentminded driver

Bounded rationality

ABSTRACT

We generalize the modified multiselves approach of Piccione and Rubinstein to (multi-player) games of imperfect recall. Four solution concepts are introduced: the multiselves agent equilibrium, the multiselves Nash equilibrium, the multiselves sequential equilibrium, and the multiselves perfect equilibrium. These modified equilibrium notions satisfy two important properties not fulfilled by the original ones. First, they always exist: every finite extensive game has at least one multiselves equilibrium of each type. Second, they form a strict hierarchy: every multiselves perfect equilibrium is a multiselves sequential equilibrium, every multiselves sequential equilibrium is a multiselves Nash equilibrium, and every multiselves Nash equilibrium is a multiselves agent equilibrium—but not conversely. Finally, in games of perfect recall, the multiselves equilibrium notions reduce to their original counterparts.

© 2018 Elsevier Inc. All rights reserved.

1. Introduction

This paper discusses and extends the classical notions of equilibrium for games of imperfect recall. Games of imperfect recall are a class of extensive games with imperfect information in which players do not always remember the experience they lived throughout the game. A player may have information at some stage of the game, and may forget that information at some later stage. This information may be about the actions of the other players (or Nature), or it may be about the player's own previous actions.

Imperfect recall is often used as a model of bounded rationality. It is employed for situations in which players have limited memory. Many real-life situations are too complex to assume individuals can have strategies that make use of the entire history of play. To deliver meaningful analyses, models of strategic interactions should account for this limitation of human behavior. A concrete illustration is the game of poker. Until recently, computer programs were performing poorly against human players. But in the past few years, programs have begun to play competitively against human players—even expert players. The breakthrough occurred when researchers started to abstract the game and compute approximate equilibria of this game. Modeling the game as one of imperfect recall has shown to deliver the best performance (Johansson et al., 2011).

In spite of their importance, games of imperfect recall are not well understood. Classical solution concepts, such as Nash equilibrium, and its refinements sequential equilibrium and perfect equilibrium, continue to be well defined. However, their

[☆] We thank the editor and two anonymous referees for helpful comments and suggestions. Lambert thanks Microsoft Research and the Cowles Foundation at Yale University for their hospitality and financial support.

* Corresponding author.

E-mail addresses: nlambert@stanford.edu (N.S. Lambert), adrian.marple@gmail.com (A. Marple), shoham@cs.stanford.edu (Y. Shoham).

use in games of imperfect recall present some challenges. With perfect recall, it is well known that mixed and behavioral strategies are equivalent: mixed strategies can be emulated by behavioral strategies, and conversely (Kuhn, 1953). Two major properties are equilibrium existence (every finite game has a Nash equilibrium, as well as a sequential equilibrium and a perfect equilibrium) and equilibrium hierarchy (every perfect equilibrium is a sequential equilibrium, which in turn is a Nash equilibrium).¹ However, with imperfect recall, the equivalence between mixed and behavioral strategies no longer hold. And with behavioral strategies—the more natural class when a game is of imperfect recall—the two key properties break down.

The fact that standard properties of equilibrium do not hold is not merely technical inconvenience. It reflects deeper conceptual unclarity, which already arise in decision problems (i.e., games with only one player) and were discussed by Piccione and Rubinstein (1997) in their provocative example “the paradox of the absentminded driver,” and the reactions that followed.² In that example, the optimal strategies of the decision maker are not time consistent, in spite of the fact that no new information arrives and the decision maker’s preferences do not change. We discuss this line of work in more details in our review of the literature below.

The issues stressed by Piccione and Rubinstein are not limited to decision problems. Imperfect recall forces us to reconsider the procedural aspects of decision making. In light of their analysis, we do not find it surprising that the two major equilibrium properties of games of perfect recall cease to hold under imperfect recall. In the classical equilibrium concepts, a player can decide of a plan of action up front and can reconsider her plan every time she takes action. But in many games of imperfect recall, it is conceptually difficult to make sense of the idea that a player can control her future actions. To resolve the absentminded driver paradox, Piccione and Rubinstein suggest that decision makers be able to reassess their immediate action at any decision point, but that they be unable to reassess their plan of future actions. This suggestion, which Piccione and Rubinstein call the *modified multiselves approach*, is a generalization of the multiselves concept of Strotz (1955). Aumann et al. (1997) and Gilboa (1997) forcefully defend this view. It also reinforces the idea that behavioral strategies are the natural class to consider in games of imperfect recall.

We find the modified multiselves approach appealing because it captures a fundamental principle of decision making—action optimality—which is complementary to the standard notions of equilibrium. Because of the commonly called “one-deviation” property, this limitation on decision making procedures does not impact outcomes in environments with perfect recall. However, it resolves the time-consistency paradox in decision problems with imperfect recall, and as we prove in this paper, it allows to recover the two major equilibrium properties aforementioned in games with imperfect recall.

Building on the modified multiselves approach, we adapt classical solution concepts to games of imperfect recall. This paper focuses on four classical solution concepts: the agent equilibrium (AE) (Strotz, 1955), the Nash equilibrium (NE) (Nash, 1951), the sequential equilibrium (SE) (Kreps and Wilson, 1982), and the perfect equilibrium (PE) (Selten, 1975). For each solution concept, we present a modification that integrates the idea of modified multiselves. We term these modified equilibrium notions, respectively, the multiselves agent equilibrium (MAE), the multiselves Nash equilibrium (MNE), the multiselves sequential equilibrium (MSE), and the multiselves perfect equilibrium (MPE).

The main result of this paper shows that every finite extensive game of imperfect recall always has a multiselves perfect equilibrium—whereas it may not have a perfect equilibrium. We also show that the usual chain of inclusion applies: every MPE is an MSE, every MSE is an MNA, every MNA is an MAE (and these inclusions are strict). Hence, in addition to having an MPE, every finite game of imperfect recall also has an MSE, MNE, and MAE. The two aforementioned technical inconsistencies are thus resolved. In addition, and importantly, the newly defined equilibrium notions agree with the original ones outside of the scope of imperfect recall. Specifically, when a game has perfect recall, we show that every MAE (resp. MNE, MSE, MPE) is an AE (resp. NE, SE, PE) and conversely. Aside from these results, we believe that our treatment provides conceptual clarity.

Naturally, the adapted solution concepts are reasonable only to the extent that one agrees with the underlying interpretation of decision making. These concepts lie at one extreme in which a player forms expectations about how she will make future decisions, but can only control her immediate action—or, as Segal (2000) notes, the player can reoptimize her plan of action at any time, but cannot commit to the new plan. An important implication of equilibrium existence is that it now becomes possible to derive predictions in any game of imperfect recall. However, there are also situations for which one may want to give the player more control over future actions. In those situations, the multiselves equilibrium notions may be less relevant.³ One may prefer (for example) to use mixed strategies paired with the traditional solution concepts—which lie at the other extreme and give full planning abilities to the player. Note that equilibrium existence continues to be guaranteed, because of the use of mixed strategies as opposed to behavioral strategies. And, of course, there are many possibilities in between, in which players can reconsider a plan of actions at some decision points but not at others (for example, as in Battigalli, 1997). In that sense, the multiselves solution concepts are complementary to the traditional ones.

The paper is organized as follows. In Section 2 we illustrate some of the main ideas in the context of a simple example. We define the four multiselves solution concepts in Section 3. In Section 4 we demonstrate that every finite game of

¹ These properties appear in the original works of Nash (1951), Selten (1975) and Kreps and Wilson (1982).

² A symposium in *Games and Economic Behavior* (1997, Vol. 20) is dedicated to discussions and reactions to the paradox.

³ The password game presented in Section 4.3 of Rubinstein (1998) is a good illustration of such a situation. The multiselves approach to decision making fails to prune away an equilibrium that, one might argue, is not realistic.

imperfect recall has at least one multiselves perfect equilibrium. In section 5 we show that the multiselves notions of equilibrium form a strict hierarchy, and thus that finite game of imperfect recall has at least one multiselves equilibrium of each of the four types. In the same section, we show that these modified solution concepts reduce to the original counterparts for games of perfect recall. Section 6 concludes.

While our paper defines four equilibrium notions, the notion of multiselves Nash equilibrium is more complex and less natural than the other three notions, despite the simplicity of its original version the Nash equilibrium. To simplify the exposition, we relegate all definitions and results related to the multiselves Nash equilibrium to Appendix A. Appendix B includes the proofs omitted from the main text.

Related literature

Imperfect recall can be seen as a model of bounded rationality, and so, from a general standpoint, our work is related to the literature on bounded rationality. Specifically, imperfect recall is a model of knowledge imperfection, in which knowledge is memory (we refer the reader to Rubinstein, 1998 for a discussion and a survey of the classical models). It relates most to two classes of bounded rationality models. In the first class, the notion of strategy is replaced by the notion of finite automaton or machine. Machine play is especially relevant in the analysis of repeated games (Neyman, 1985; Rubinstein, 1986; Abreu and Rubinstein, 1988; Papadimitriou and Yannakakis, 1994). Because a finite automaton is limited in complexity by the number of its states, the amount of memory and knowledge a player can accumulate throughout the game is bounded—as it can be for a player who has imperfect recall. However, with imperfect recall, there no limitation of the set of strategies a player can choose from, these strategies can be arbitrarily complex. The imperfections are focused on what players know, or the sort of inferences players can make, and are specified as part of the description of the game tree. The second class models limited foresight. Players with limited foresight have difficulties understanding the extensive game they confront, which describes the future rounds of play. In contrast, with imperfect recall, players have a clear understanding of the game tree and the consequences of their own future actions and those of others, but have difficulties remembering past rounds of play. For a discussion of limited-foresight reasoning and its implications, see, for example, Jehiel (1995, 2001), Samuelson (2001) and Wichardt (2010). Another interpretation of imperfect recall is that of players representing organizations which consist of several agents who share a common preference but have different roles and lack the ability to coordinate or communicate their actions (Isbell, 1957). Because our solution concepts rely on the idea of optimality at the action level, they prevent coordination among different agents and so are particularly relevant to this interpretation.

In substance, our paper is related to the stream of the literature that illustrates the subtleties of imperfect recall. In particular, Wichardt (2008) shows that Nash equilibria do not always exist in games with imperfect recall.⁴ Similarly, Kline (2005) shows that the usual chain of inclusion breaks down, in that some sequential equilibria are not Nash equilibria and some perfect equilibria are not sequential equilibria. The fact that the key properties of games of perfect recall no longer hold raises the question of whether the classical solution concepts are well adapted to environments with imperfect recall, it is the starting point of our work. In contemporaneous work, Halpern and Pass (2016) and Hillas and Kvasov (2017) also argue that the classical solution concepts are inappropriate, and propose alternative equilibrium notions. Halpern and Pass define ex ante notions of sequential and perfect equilibrium. They require that a player's strategy be a best response at any information set (or collection of information sets) among strategies that do not affect the player's utility at histories not coming after the information set (or collection) being considered. Hillas and Kvasov extend the notions of sequential and quasi-perfect equilibrium, with the objective of recovering the central ideas of backward induction in a way compatible with imperfect recall. Importantly, both papers work with the notion of mixed behavioral strategies, which encompass both mixed strategies and behavioral strategies, they define generalized beliefs that span multiple information sets, and allow for the player to deviate simultaneously at several information sets. Because these works allow for mixed strategies, equilibrium existence, which already holds for the classical solution concepts, is not their main result. In contrast, our focus is on behavioral strategies and action optimality, for which equilibrium existence becomes a more central contribution. To finish, we should mention the idea of distributed games defined by Monderer and Tennenholtz (1999). In a distributed game a given standard game is played in parallel at different "locations," with messages broadcast from a given location to the others when the game at that location has been played. Although our work is not directly connected, both works utilize an idea of parallelism and separation of decisions.

An earlier literature that originates with the work of Piccione and Rubinstein (1997) illustrates conceptual difficulties that arise in games of imperfect recall. Their main example features a driver who engages in a highway and may take one of two exits. Taking the first exit generates payoff 0, while taking the second exit generates payoff 4. The driver may also choose not to exit at all, which yields payoff 1. The game is reproduced in Fig. 1. The particularity of the example is that the driver is not able to tell which exit she is facing. Piccione and Rubinstein show that the driver's optimal strategy is not time consistent: the optimal strategy at the entry of the game is no longer optimal once the decision maker faces the choice either to exit or to continue. The example seemingly violates common sense because the decision maker of this game receives no new information as she moves down the game tree, and her preferences are fixed once and for all. Hence its paradoxical nature. It has sparked a number of reactions and possible resolutions (see, in particular, Aumann et al. 1997; Battigalli 1997; Gilboa 1997; Grove and Halpern 1997; Halpern 1997; Segal 2000).

⁴ The nonexistence has also been noted informally by Başar and Olsder (1999) in their Example 2.4.




Fig. 1. The absentminded driver.

The origin of the paradox lies in the interpretation of how the decision maker can control her behavior. Aumann et al. (1997) argue that at any point in time, the decision maker should only be able to control her current action, but not her future actions. This assumption is what Piccione and Rubinstein refer to as the modified multiselves approach, which they propose as a possible interpretation to address the paradox. We agree with this interpretation, which forms the basis of our modified equilibrium concepts. In particular, our notion of multiselves agent equilibrium is a generalization of the modified multiselves approach to multiplayer environments. Relatedly, Battigalli (1997) defines the modified multiselves sequential equilibrium for decision problems, and our notion of multiselves sequential equilibrium is the generalization to multiple players.

The proof of our main result appeals to a reformulation of the original game of imperfect recall that we term “multiselves agent form.” This reformulation generalizes the agent form of a game defined by Kuhn (1953), and is directly inspired by the comment of Gilboa (1997). Gilboa offers a more ideological treatment to the absentminded driver example. He argues that the paradox arises because of the way Piccione and Rubinstein model the decision problem. According to Gilboa, “decision problems can and should be formulated in such a way that information sets do not contain more than one decision node on each path.” Gilboa proposes to formulate the decision problem as a game of perfect recall, in which the decision maker is replaced by a collection of two agents who make choices independently on behalf of the decision maker, taking one another’s behavior as given. As Lipman (1997) outlines, although different in form, the comments of Gilboa and those of Aumann, Hart and Perry agree on substance. The multiselves agent form of this paper operates a similar transformation of the original game, which enables us to leverage the simplicity of perfect recall. This transformed game also stratifies information sets, and to do so, makes use a notion of frontier of information sets that is related to the notion of upper frontier by Halpern (1997).

This stream of literature focuses on decision problems and, for a large part, works with specific examples. The key message is that with the proper interpretation of the decision maker’s behavior, there are no inconsistencies. This observation applies beyond decision problems. Classical equilibrium notions are not well adapted to environments with imperfect recall, because implicit in those notions is the idea that a player is able to reassess her plan of actions at the level of information sets. However, the ideas developed in this literature did not culminate in a map of solution concepts for general multiplayer games that naturally extends the map from the perfect recall case. One purpose of this paper is to fill this gap.

2. A motivational example

In this section we present some of the main ideas of this paper in the context of a simple example. We consider a two-player extension of Piccione and Rubinstein’s absentminded driver, that we refer to as the *absentminded driver and the policeman*. The game is shown in Fig. 2, and is used throughout the paper to exemplify the concepts.

In this game variation of the baseline decision problem, the driver (who continues to be a female) has committed a misdemeanor and is being chased by a policeman (a male). As in the original decision problem, the driver is on the highway and may choose either to exit at one of the two possible exits, or to continue. The policeman must anticipate the driver’s choice in order to catch her. He can choose to wait on the highway (action *H* on the game tree) or to wait at the first exit (action *E*). A camera is located at the second exit, so that if the driver takes that exit, she is caught by the policeman for sure. If the driver is caught by the policeman, the payoffs are 0 for both players. If the policeman fails to catch the driver because he waits at the wrong location, the payoff for the driver is 1 and the payoff for the policeman is -1 . The game is a simple variant of the standard zero-sum inspection game. Note that, as in the original absentminded driver decision problem, the driver is not able to identify the exit she faces when she must decide whether to exit or to continue.

This game does not have a Nash equilibrium (in behavioral strategies). The argument is as follows. Suppose the driver chooses to continue on the highway with probability p (no matter which exit she faces, as she cannot distinguish between the two), and the policeman chooses to wait for the driver at the first exit with probability q . With this strategy profile,




Fig. 2. The game of the absentminded driver and the policeman.




Fig. 3. An extension of the absentminded driver and the policeman.

the driver never exits with probability p^2 , and leaves the highway at the first exit with probability $1 - p$. Thus, the overall payoff of the driver is $p^2 \times q + (1 - p) \times (1 - q)$, which is a strictly convex function in p . Therefore, mixing yields a strictly worse payoff than the best pure strategy. And, if in all Nash equilibria, the driver always chooses a pure strategy, finding an equilibrium amounts to finding a pure strategy Nash equilibrium in the game of matching pennies, where the driver wants to mismatch and the policeman wants to match. Of course, such equilibrium does not exist.

Even for games that have a Nash equilibrium, the equilibrium refinements that are the sequential and perfect equilibrium—needed to prune away the nondesirable Nash equilibrium—do not necessarily exist. To illustrate, consider the simple extension of the game presented in Fig. 3. In this extended game, both driver and policeman have additional choices to make before the driver enters the highway. As the chase begins, the driver can decide to surrender immediately to the policeman (action S). In this case, the payoff is zero for both players. Or she can speed up and try to escape (action T). If she tries to escape, the policeman can also speed up and intercept the driver's vehicle before she enters the highway (action I), however, the interception causes a car accident and both players get payoff -1 . Or, the policeman can wait until the driver gets on the highway and avoid the accident (action W).

This game has a Nash equilibrium, in which the driver surrenders immediately and the policeman is prepared to intercept the driver if the driver does not surrender immediately. However, this equilibrium involves a noncredible threat. In games of perfect recall, using the sequential and perfect equilibrium notions would exclude such strategic behavior. However, this particular game neither has a sequential equilibrium nor a perfect equilibrium, for a similar reason as the game of Fig. 2 not having a Nash equilibrium.

These two examples motivate the need for alternative equilibrium notions suitable to environments with imperfect recall. What causes nonexistence of classical equilibria in these examples is no different from what causes the time inconsistencies in the absentminded driver decision problem. The origin of the problem is that when the driver, on the highway, contemplates a possible deviation from any conjectured equilibrium strategy, the deviation is at the level of the information set, which implies not only a change of the immediate action, but also a change of future actions when she is at the level of the first intersection.

To remedy this situation, let us assume that, as the driver contemplates a deviation, she contemplates a deviation only once, at the instance at which she operates. So, if after making a decision at her information set she visits the same information set a second time, she anticipates she will follow the strategy originally planned, as opposed to any new action plan. This idea is central in the modified multiselves approach of Piccione and Rubinstein. This sort of “one-shot deviation” is captured by what we call a “phantom strategy.” Phantom strategies are defined formally in the next section and are the main primitive behind all our multiselves solution concepts.

Let us focus on the game of Fig. 2. Consider the strategy profile in which, with probability $p = 2/(\sqrt{5} + 1) \approx 62\%$, the driver continues on the highway (as opposed to exiting), and with probability $q = 1/\sqrt{5} \approx 45\%$, the policeman tries to catch the driver at the first exit. We now proceed to show that this strategy profile is an equilibrium, in the sense that each player best responds to the other, accounting for the “one-shot deviation” restriction from the original strategies just mentioned.

To properly evaluate the impact of changing her action, the driver must form beliefs about what exit she faces. If she follows the prescribed conjectured equilibrium strategy, the driver finds herself at the level of the first exit $1/(1 + p) \approx 62\%$ fraction of the time, and at the level of the second exit $p/(1 + p) \approx 38\%$ fraction of the time. From the driver’s viewpoint, these numbers correspond to the respective probabilities of facing the first and second exit at the time she must take action. These probabilities form the driver’s *beliefs*.

If she is at the level of the second exit, then her subjective expected payoff is $p \times q$ when following the strategy prescribed by the equilibrium conjecture. If she is at the level of the first exit, then her subjective expected payoff is $p \times pq + (1 - p) \times (1 - q)$.

Can she do better? Suppose she chooses to continue on the highway with probability p' instead of p at the time she is called upon to take action—wherever she is, as she does not know where she is. If she happens to be at the first intersection, she gets payoff $p' \times pq + (1 - p') \times (1 - q)$. If she happens to be at the second intersection, she gets payoff $p' \times q$. Overall, accounting for the uncertainty about which exit she faces, her overall subjective expected payoff from deviating is

$$\frac{1}{1 + p} \times (p' \times pq + (1 - p') \times (1 - q)) + \frac{p}{1 + p} \times p' \times q = \frac{3}{\sqrt{5}} - 1,$$

which does not depend on p' . Hence, the choice of $p' = p$ is optimal for the driver, considering as fixed the policeman’s strategy.

Now let us verify that the policeman’s strategy is also optimal. This case is easier, because the policeman is not absentminded: in any game instance, the policeman plays in his information set only once. According to her conjectured equilibrium strategy, the driver remains on the highway with probability p^2 , and leaves the highway at the first exit with probability $1 - p$. Thus, if the policeman chooses to wait at this first exit with probability q' , he gets overall expected payoff

$$-p^2 \times (1 - q') - (1 - p) \times q' = -\frac{2}{\sqrt{5} + 3}$$

which does not depend on q' . Hence, $q' = q$ is a best response.

Overall, if the criterion for optimality is that all strategies resist the “one-shot deviation” explained above, then the strategy profile (p, q) has both players mutually best respond to each other, and so is, in that sense, an equilibrium.

This section presents a basic example and its arguments remain heuristic and incomplete. In the next section, we define precisely the notion of phantom strategy that captures the idea of one-shot deviation in the modified multiselves approach. We use the notion to define the four equilibrium notions (one of which is relegated to Appendix A). In the language of the next section, the equilibrium found above is a multiselves perfect equilibrium, which is the most refined of the four solution concepts—and so it is also a multiselves sequential, Nash, and agent equilibrium, by the inclusion property shown in Section 5. The example displayed here generalizes substantially, and this generalization is the main purpose of this paper. In Section 4, we demonstrate that every finite extensive game of imperfect recall has a multiselves perfect equilibrium.

3. Multiselves solution concepts

3.1. Preliminary definitions

We begin with several standard definitions and notations. Throughout, all games are assumed finite. The focus is on extensive games and in the sequel, we refer to a (finite) extensive game simply as a “game,” unless mentioned otherwise. As common in the literature, each game is represented as a tuple $\Gamma = \langle N, H, P, \rho, u, \mathcal{I} \rangle$. We assume the reader is familiar with this representation (see, for example, Chapter 11 of Osborne and Rubinstein (1994) for full details) and we briefly review the notation as follows.

- N is the finite set of players.
- H is the finite collection of the possible histories of actions. We denote by Z the set of terminal histories, and by $A(h)$ the set of actions available at history h .
- $P: H \setminus Z \rightarrow N \cup \{c\}$ is the player function, it assigns a player of N or Nature (player “c”, for “chance node”) to every nonterminal history.

- ρ is a prior on the actions of Nature: if Nature is to play at history h , $\rho(h)$ denotes the probability distribution over the possible actions to be played.
- $u_i: Z \rightarrow \mathbf{R}$ is the utility function that captures the preferences of player i . It records the player's payoff for every terminal history. The player then evaluates lotteries over game outcomes by their expected utility, and so behaves as an expected utility maximizer.
- \mathcal{I} is the collection of all the information sets of the game. It partitions the collection of nonterminal histories. Because, in any information set, the available actions and the player who operates are required to be the same for all histories in that information set, we abuse notation and write, for information set I , $P(I)$ the player assigned to play at I , and $A(I)$ her set of available actions. We denote by \mathcal{I}_i the collection of the information sets at which player i operates.

The *experience* of player i at history h records what the player “lived” along history h . It is defined as the sequence of information sets and actions of that player along the history h and is denoted $\text{exp}_i(h)$. A game has *perfect recall* if for every information set I , the player assigned to play at I has the same experience at all histories of I . Otherwise, the game has *imperfect recall*. Concretely, a game with imperfect recall is a game in which at least one player, at some history at which she operates, is unsure about how she got to that point, because she believes that two or more different experiences can explain where she is. Informally, imperfect recall captures a player's limited memory.

A history $h = (a_1, \dots, a_K)$ precedes history $h' = (a_1, \dots, a_L)$ when $K \leq L$ (written $h \leq h'$), and h strictly precedes h' if $K < L$ (written $h < h'$). A game has *absentmindedness* if there exist two histories of a same information set such that one history strictly precedes the other. In a game with absentmindedness, some player can visit the same information set more than once as the game is being played. As she takes actions, she does not remember whether she has already visited the information set.

A *pure strategy* for a player assigns a single valid action to every information set of that player, and a *mixed strategy* is a distribution over the player's pure strategies. In contrast, a *behavioral strategy* assigns a distribution over actions to every information set of the player. Behavioral strategies are the more natural class to consider in games of imperfect recall and are the focus on this paper.⁵ Unless mentioned otherwise, in the sequel we refer to behavioral strategies simply as “strategies.” The set of possible (behavioral) strategies for player i is denoted Σ_i , and the set of all possible (behavioral) strategy profiles for all players is denoted Σ .

For a given strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ and an information set I , $\sigma(I)$ denotes the distribution over actions induced by σ at I (i.e., $\sigma(I) = \sigma_i(I)$ where i is the player assigned to play at I). So, given a particular action a that can be played at I , $\sigma(I)(a)$ is the probability that a will be played upon reaching I . A strategy profile σ also induces a probability distribution over terminal histories. This distribution is denoted π_σ . We extend π_σ to be defined at all histories of the game by $\pi_\sigma(h) = \sum_{z \in Z: h \leq z} \pi_\sigma(z)$. The induced distribution over terminal histories enables us to conveniently write the expected utility of player i when all the players follow strategy profile σ ,

$$u_i(\sigma) = \sum_{z \in Z} \pi_\sigma(z) \cdot u_i(z).$$

We now turn to the definitions of the classical solution concepts that we generalize, and formalize the notion of consistent beliefs. Consistent beliefs will also be used in the definition of the multiselves solution concepts.

Nash equilibrium A strategy profile σ is a *Nash equilibrium* (NE) if, for every player i , and every strategy σ'_i of that player, $u_i(\sigma'_i, \sigma_{-i}) \leq u_i(\sigma)$.

Consistent beliefs and sequential equilibrium In a sequential equilibrium, players form subjective beliefs about where they are in the game, and best respond accordingly. For every information set I , a belief $\mu(I)$ at that information set is a distribution over the histories of I . It represents a player's belief about the likelihood of operating at a node of I versus another, conditionally on being at I . The belief μ is *consistent* with a strategy profile σ if for every information set I that is reached with positive probability under σ , and every history $h \in I$, we have

$$\mu(I)(h) = \frac{\pi_\sigma(h)}{\sum_{h' \in I} \pi_\sigma(h')}. \quad (1)$$

This definition of consistency is the one retained by Piccione and Rubinstein (1997). The usual interpretation is that beliefs should “satisfy Bayes rule whenever possible.” However, as Piccione and Rubinstein observe, Equation (1) is not innocuous for games of imperfect recall, because there may be instances of the game in which a player finds herself in the same information set more than once. In this case, two different histories of the same information set should no longer be viewed as incompatible events (in the probability theory sense) and the probabilities of these two events should not simply be summed to get the overall probability.

⁵ Note that under mixed strategies, the two key properties discussed in the Introduction, existence and equilibrium hierarchy, continue to hold in games of imperfect recall.

Simply put, having beliefs satisfy Equation (1) can be interpreted as having beliefs that are “frequency probabilities.” If the game was being played repeatedly infinitely often, then $\mu(I)(h)$ would equal the fraction of times that the player, who plays at information set I , is at history h , conditionally on her being at I . We find this long-run frequency approach to belief formation intuitive and reasonable. However, other approaches exist that define consistent beliefs differently. For further discussion on the subject, we refer the reader to Grove and Halpern (1997), Halpern (1997), and Piccione and Rubinstein (1997). It is worth noting that the definition of consistent beliefs as Equation (1) is substantial in our existence results of Section 4. Whether existence holds under the alternative definitions of consistency remains an open question, and would necessitate a different way of proof.

Similarly to π_σ , we can define the distribution over terminal histories induced by strategy profile σ conditional on the play starting from history h , $\pi_\sigma(\cdot | h)$. The expected utility of player i conditional on being at history h is then $u_i(\sigma | h)$, which can be written

$$u_i(\sigma | h) = \sum_{z \in Z} \pi_\sigma(z | h) \cdot u_i(z).$$

Combining subjective beliefs and conditional utilities, we get a subjective version of a player’s utility. For a given strategy profile σ , the subjective utility of player i at information set I is defined as

$$SU_i(\sigma; I, \mu) = \sum_{h \in I} \mu(I)(h) \cdot u_i(\sigma | h).$$

It reflects the payoff of that player when she is playing at I , and when her beliefs about which node she is at within I are specified by distribution $\mu(I)$.

A strategy profile σ is a *sequential equilibrium* (SE) if there exists a sequence of completely mixed strategy profiles $\sigma^1, \sigma^2, \dots$ that converges to σ , and a sequence of positive reals $\varepsilon_1, \varepsilon_2, \dots$ that converge to 0, such that for every k , for every belief μ consistent with σ^k , for every player i and every information set I assigned to player i , and for every strategy σ'_i ,⁶

$$SU_i((\sigma'_i, \sigma_{-i}^k); I, \mu) \leq SU_i(\sigma^k; I, \mu) + \varepsilon_k.$$

Perfect equilibrium The notion of perfect equilibrium is based on the notion of perturbed games. Given a game Γ , a *perturbation* η of Γ is a function that assigns to every information set I and every action available at I a positive probability, such that, for any given information set, the sum of the probabilities over actions is always less than one. A *perturbed game* is a pair (Γ, η) . In a perturbed game, players are not allowed to use strategies that play some action with less than η probability, i.e., $\sigma_i(I)(a) \geq \eta(I)(a)$ for every player i , every information set I at which i operates, and any action a available at I . These are the *perturbed strategies* of (Γ, η) .

Given a game Γ , strategy profile σ is a *perfect equilibrium* (PE) if there exists a sequence of perturbed games, $(\Gamma, \eta_1), (\Gamma, \eta_2), \dots$, with $\eta_k \rightarrow 0$, and a sequence of correspondingly perturbed strategy profiles $\sigma^1, \sigma^2, \dots$ that converges to σ , such that for every k , σ^k is a Nash equilibrium of the game (Γ, η_k) .

Agent equilibrium Introduced by Kuhn (1953), the concept of agent equilibrium is perhaps not as well known as the other three, but it plays an important role in games of imperfect recall and in our definitions of the multiselves solution concepts.

Kuhn defines the agent form of an extensive game as a modified version of the game that gives control of every information set to an independent agent who receives the same payoff as the player who operates at that information set in the original game. Intuitively, the agent form of a game delegates a player’s decisions to several independent agents, one agent per information set. An *agent equilibrium* (AE) of game Γ is a strategy profile that is a Nash equilibrium of the agent form of Γ . Our multiselves solution concepts are related to the agent form, because in these equilibria, players’ decisions are delegated to even finer entities that break up the information sets.

3.2. Phantom strategies

The multiselves equilibrium notions rest on an hypothetical construct we term *phantom strategies*. A phantom strategy represents a possible deviation from equilibrium a player may contemplate.

Whereas a standard strategy assigns a distribution over actions to every information set of a player, a phantom strategy specifies a distribution over actions at *every history at which a player operates*, thus ignoring information sets. Precisely, a phantom strategy for player i , σ_i , assigns a distribution $\sigma_i(h)$ over the actions that player i may take at history h .

Phantom strategies enable us to represent the reasoning of a player who confronts a choice at a some given information set. Of special interest is a subclass of phantom strategies we call *single-deviation phantom strategies*. For a given player i playing at information set I , any history h in that set, and a probability distribution p over actions available to that player at I , the single-deviation phantom strategy $\sigma_i[h/p]$ is defined as

⁶ Kreps and Wilson (1982) provide another common characterization of sequential equilibrium in terms of what they term *strongly consistent beliefs*.

$$\sigma_i[h/p](h') = \begin{cases} p & \text{if } h' = h, \\ \sigma_i(I) & \text{otherwise.} \end{cases}$$

Note that, since a history uniquely defines the information set that includes it, we need not specify I as part of the phantom strategy $\sigma_i[h/p]$. In some instances it is useful to consider deviations at more than one history. If H is a nonempty set of histories of information set I , we define the phantom strategy $\sigma_i[H/p]$ as

$$\sigma_i[H/p](h') = \begin{cases} p & \text{if } h' \in H, \\ \sigma_i(I) & \text{otherwise.} \end{cases}$$

The special case $H = I$ is at the opposite of the single-deviation strategy, it corresponds to a *uniform-deviation strategy*. As for regular strategies, a profile of phantom strategies (single deviation or not) induces a distribution over terminal histories (both from an ex ante perspective and conditionally on some given history of play), which can be extended to all histories—the definitions are analogous to the case of regular strategies and are omitted. In particular, the notion of player utility extends directly to profiles of phantom strategies (and so to mixed profiles).

The single-deviation strategy $\sigma_i[h/p]$ reflects the idea of one-shot deviation: it can be interpreted as player i about to play at history h , and deviating from her prescribed strategy just once at this particular instance of decision, by choosing an action according to distribution p , and following the originally prescribed strategy σ_i for all future decisions. Of course, unless the information set is a singleton, the player is generally unable tell with confidence which history she is playing from. This is the reason why we call these strategies “phantom” strategies. It remains an hypothetical deviation, *as if* the player *knew* she was at a particular history within his information set.

However, even if the player does not know which history of an information set she is playing from, she is still able to form beliefs about the likelihood of being at one history versus another. This belief allows her to evaluate the impact of a one-shot deviation from the prescribed strategy at any given information set. The evaluation is captured by what we call *multiselves subjective utility*. Given a strategy profile σ , the multiselves subjective utility of a player i who operates at information set I and deviates just once from her original strategy by choosing an action distributed according to p is

$$\text{MSU}_i(\sigma; I, p, \mu) = \sum_{h \in I} \mu(I)(h) \cdot u_i((\sigma_i[h/p], \sigma_{-i}) | h),$$

where μ denotes the belief of that player—which, in particular, characterizes her personal assessment about where in the information set she is playing from. Of course, when there is no deviation from the original strategy, the multiselves subjective utility agrees with the subjective utility: $\text{SU}_i(\sigma; I, \mu) = \text{MSU}_i(\sigma; I, \sigma(I), \mu)$.

3.3. Multiselves equilibrium notions

We can now define precisely the multiselves analog for each of the four solution concepts presented in Section 3.1. Note that these new concepts are motivated by games of imperfect recall, but continue to be well defined for games of perfect recall.

We begin with the multiselves agent equilibrium, which is instrumental in our results and in the definition of the other solution concepts.

Definition 1. A strategy profile σ is a *multiselves agent equilibrium* (MAE) if there exists a belief μ consistent with σ such that, for every player i , for every information set I assigned to player i and reached with positive probability, and for every probability distribution p over actions available at I , $\text{MSU}_i(\sigma; I, p, \mu) \leq \text{SU}_i(\sigma; I, \mu)$.

The multiselves agent equilibrium is to multiplayer games what Piccione and Rubinstein’s modified multiselves approach is to decision problems. It captures the “one-shot deviation” criterion of optimality: when re-evaluating her strategy at any instance at which she operates, a player can reassess her immediate action, but not her plan of future actions. For concreteness, we illustrate this concept on Piccione and Rubinstein’s absentminded driver problem and our simple game of Section 2 in the two examples below.

Example 1. Consider the one-player game of Section 1 presented in Fig. 1. Let us start by computing the ex ante utility-maximizing strategy for the driver. If p denotes the probability of choosing action continue (the driver’s strategy), then the driver’s expected utility is $u(p) = p^2 + 4p(1 - p)$, whose maximum is reached at $p = 2/3$. Piccione and Rubinstein showed that any ex ante optimum is “modified multiselves time consistent,” meaning that it remains optimal at all stages of the game under the modified multiselves approach. Since our notion of multiselves agent equilibrium is a multiplayer generalization of the modified multiselves approach, one would expect that strategy $p = 2/3$ is a multiselves agent equilibrium. Let us confirm this intuition.

First, we note that the driver can only have one belief consistent with the prescribed strategy—namely, the belief that she is at the first intersection with probability $3/5$, and at the second intersection with probability $2/5$. There is only one

information set at which the driver can re-evaluate her strategy. The multiselves subjective utility following the one-shot deviation p' at that information set is

$$\frac{3}{5} \times p'(p + 4(1 - p)) + \frac{2}{5} \times (p' + 4(1 - p')) = \frac{8}{5}.$$

Since this utility does not depend on p' , the decision maker cannot improve on her strategy of $p = 2/3$. Hence, strategy $p = 2/3$ is a multiselves agent equilibrium.

Example 2. We now return to our leading example of Section 2 presented in Fig. 2. The driver's strategy is captured by the probability p of the driver choosing to continue on the highway (action C), and the policeman's strategy is captured by the probability q of trying to catch the driver at the first exit (action E). We proceed to show that the profile of strategies $p = 2/(\sqrt{5} + 1)$ and $q = 1/\sqrt{5}$ is a multiselves agent equilibrium, this time using the formal tools developed in this section.

Let h_1 be the history at the first intersection (the null history) and let h_2 be the history at the second intersection (history C). The driver (player 1), has only one information set, denoted I_1 , composed of h_1 and h_2 . Following notation of Section 3.2, $p[h_1/p']$ is the single-deviation phantom strategy according to which the driver continues with probability p' at h_1 , and $p[h_2/p']$ is the analog at h_2 . The driver's utilities from using these phantom strategies at the two intersections are respectively

$$u_1((p[h_1/p'], q) | h_1) = p' \times pq + (1 - p') \times (1 - q),$$

$$u_1((p[h_2/p'], q) | h_2) = p' \times q.$$

The only belief consistent with strategy profile (p, q) assigns respective probabilities $\mu(I)(h_1) = 1/(1 + p)$ and $\mu(I)(h_2) = p/(1 + p)$ to histories h_1 and h_2 conditionally on being at I_1 . Therefore, the driver's multiselves subjective utility from one-shot deviation p' at I_1 is

$$\begin{aligned} \text{MSU}((p, q); I_1, p', \mu) &= \mu(I)(h_1) \cdot u_1((p[h_1/p'], q) | h_1) + \mu(I)(h_2) \cdot u_1((p[h_2/p'], q) | h_2) \\ &= \frac{1}{1 + p} \cdot (p' \times pq + (1 - p') \times (1 - q)) + \frac{p}{1 + p} \cdot (p' \times q) \\ &= \frac{3}{\sqrt{5}} - 1. \end{aligned}$$

As its value is independent of the deviation p' , the driver cannot improve on her original strategy p .

The policeman also has only one information set I_2 , which includes the history $h_3 = (C, C)$ (the policeman awaits on the highway) and history $h_4 = (E)$ (the policeman awaits at the first exit). The policeman's utilities from the single-deviation phantom strategies that deviate from q to q' at histories h_3 and h_4 respectively are

$$u_2((p, q[h_3/q']) | h_3) = -q',$$

$$u_2((p, q[h_4/q']) | h_4) = -(1 - q').$$

The only belief μ that is consistent with profile (p, q) satisfies $\mu(I_2)(h_3) = p^2/(1 - p + p^2)$ and $\mu(I_2)(h_4) = (1 - p)/(1 - p + p^2)$. Therefore, the multiselves subjective utility of the policeman from one-shot deviation q' at I_2 is

$$\begin{aligned} \text{MSU}_2((p, q); I_2, q', \mu) &= \mu(I_2)(h_3) \cdot u_2((p, q[h_3/q']) | h_3) + \mu(I_2)(h_4) \cdot u_2((p, q[h_4/q']) | h_4) \\ &= \frac{p^2}{1 - p + p^2} \cdot (-q') + \frac{1 - p}{1 - p + p^2} \cdot (q' - 1) \\ &= -\frac{1}{2}. \end{aligned}$$

As its value is independent of q' , the policeman cannot improve on her original strategy q . Since neither player can improve her payoff by a one-shot deviation, the strategy profile (p, q) is a multiselves agent equilibrium.

A multiselves sequential equilibrium is defined similarly to a multiselves agent equilibrium, except for the fact that we no longer ignore the information sets that are reached with zero probability.

Definition 2. A strategy profile σ is a *multiselves sequential equilibrium* (MSE) if there exists a sequence of completely mixed strategy profiles $\sigma^1, \sigma^2, \dots$ that converges to σ and a sequence of positive reals $\varepsilon_1, \varepsilon_2, \dots$ that converges to 0 such that for every k , for every belief μ consistent with σ^k , for every player i and every information set I assigned to player i , and for every distribution p over the actions available at I , $\text{MSU}_i(\sigma^k; I, p, \mu) \leq \text{SU}_i(\sigma^k; I, \mu) + \varepsilon_k$.

The definition of multiselves perfect equilibrium is a straightforward adaptation of the original perfect equilibrium, except that the multiselves agent equilibrium is used in place of the Nash equilibrium.

Definition 3. Given a game Γ , strategy profile σ is a *multiselves perfect equilibrium* (MPE) if there exists a sequence of perturbed games $(\Gamma, \eta_1), (\Gamma, \eta_2), \dots$ for which $\eta_k \rightarrow 0$ and a sequence of correspondingly perturbed strategy profiles $\sigma^1, \sigma^2, \dots$ that converges to σ such that σ^k is a multiselves agent equilibrium of (Γ, η_k) for every k .

As opposed to the other three solution concepts, we did not find a natural multiselves analog of the Nash equilibrium well suited to imperfect recall. This is because in a Nash equilibrium, players always decide and evaluate their plan of action before entering the game, a behavior which conflicts with the decentralized decision making inherent to the multiselves approach. Consequently our definition of multiselves Nash equilibrium is somewhat contrived. It “works,” but we do not find it natural, and we suspect that no natural definition exists. For completeness we include a definition and prove the properties that it satisfies as we do for the other solution concepts, but relegate the definition and results to Appendix A.

As their original counterparts, the last two equilibrium notions are more suitable to extensive games than the multiselves agent equilibrium because they eliminate the possibility for noncredible threats.

Example 3. Consider the game depicted in Fig. 3 in Section 2. Any strategy profile that incorporates the threat of the policeman choosing to intercept the driver (and the driver surrendering to the policeman) is both a Nash equilibrium and also a multiselves agent equilibrium. Standard intuition tells us that the policeman’s threat is not credible, and the concepts used to resolve this sort of games are not Nash but sequential and perfect equilibrium, which, in this game, do not exist. However, the game has both a multiselves sequential and perfect equilibrium, in which the driver speeds up and the policeman chooses not to intercept the driver, and once the driver has entered the highway, both players behave as in the equilibrium of Example 2.

4. Main result: equilibrium existence

In this section, we prove the following result.

Theorem 1. *All games have a multiselves perfect equilibrium.*

In the next section we show that Theorem 1 implies existence of the other types of multiselves equilibrium.

The proof relies on a reformulation of the original game that we term *multiselves agent form*. Inspired by Gilboa (1997), it is a generalization of the agent form of Kuhn (1953) used in the multiselves approach advocated by Strotz (1955), in which the information sets are subdivided at finer levels. The key benefit of this transformation is to end with a game of perfect recall, which is more tractable. The proof proceeds in three steps.

1. Starting from any game, we create its multiselves agent form.
2. We show that when strategy profiles are required to satisfy a particular symmetry property, there is a one-to-one mapping between the strategy profiles of the multiselves agent form and the strategy profiles of the original game. Importantly, the connection also applies to equilibria: we show that a profile is a perfect equilibrium of the original game if and only if its multiselves agent form analog is a symmetric perfect equilibrium of the transformed game, in a sense made precise in Section 4.2.
3. We show that the resulting game, which is a standard extensive game with perfect recall, has a perfect equilibrium which satisfies the said symmetry property.

The remainder of this section elaborates on each step and proves Theorem 1.

4.1. Multiselves agent form

The multiselves agent form of a game (of imperfect recall) is a reformulation of that game, in which each player is divided into several agents, and each agent acts on behalf on the player but independently of the other agents. It embeds naturally the idea of one-shot deviation by delegating the decision power of the player to several independent agents. The main benefit of this reformulated game is to have perfect recall, whose properties are well understood.

Of course, there are many ways to “divide the player,” that is, there are multiple ways to allocate the histories at which a player operates to the different agents who represent that player. For example, in the standard agent form of a game (Kuhn, 1953), each agent controls one full information set of the player. However, for our purpose this division is not sufficient, because when the original game is of imperfect recall, the resulting agent form remains a game of imperfect recall. For example, the game of the absentminded driver and the policeman game presented in Section 2 are their own agent form.

To obtain perfect recall, we use a finer allocation in which information sets are partitioned and stratified into subsets that have no absentmindedness. The idea is inspired from Gilboa (1997), who explains that the situation of the absentminded

driver of Piccione and Rubinstein can (and should, according to Gilboa) be interpreted a game of perfect recall by viewing the driver as a collection of two agents, playing one after the other in a random order. Our definition of multiselves agent form could be read as a formalization and a generalization of Gilboa’s idea.

In the multiselves agent form, different agents operate at different information sets, but several agents may be assigned to play at the same information set on different subsets of histories. These subsets correspond to the frontiers of the histories of different depths, which extend the notion of upper frontier proposed by Grove and Halpern (1997).

We denote by $d_I(h)$ the *depth of history* h within information set I , defined recursively as follows:

$$d_I(h) = \begin{cases} 0 & \text{if } h \notin I, \\ 1 & \text{if } h \in I, \text{ and for all } h' < h, h' \notin I, \\ 1 + \max_{h' < h, h' \in I} d_I(h') & \text{otherwise.} \end{cases}$$

The depth of history h in some information set indicates how many times the player has visited that information set upon reaching h . The k -th frontier of an information set I is denoted \widehat{I}_k , and defined as the set of all histories of I that have depth k . We call depth of information set I and denote by $D(I)$ the maximum depth over the histories of I . For example, in the game of the absentminded driver and the policeman of Section 2, the null history (first intersection) has depth 1 in the driver’s information set, and history (C) (second intersection) has depth 2, so that the depth of the driver’s information set is 2. In contrast, all histories of the information set of the policeman have depth 1, and so the depth of the policeman’s information set is 1. The reason why the depth of the two information sets differ is that the driver is absentminded whereas the policeman is not. Information sets of a player that do not have absentmindedness always have depth 1.

In the agent form, each information set is controlled by an independent agent. In the multiselves agent form each frontier is controlled by an independent agent. In addition, to represent the idea that, as in the original game, an absentminded player does not know with precision where he operates inside the information set, we introduce a randomization over the agents of the transformed game: Nature chooses the agents who control the different frontiers of an information set uniformly at random, and without revealing to them the result of this randomization. The formal definition is given below, and the intuition behind it follows. It uses the game notation presented in Section 3.1.

Definition 4. The *multiselves agent form* of game Γ is a game $\widetilde{\Gamma}$ in which the players are called agents and defined as follows.

- The set of agents \widetilde{N} is the set $\bigcup_{I \in \mathcal{I}} \{I\} \times \{1, \dots, D(I)\}$. Let $\widetilde{n} = |\widetilde{N}|$ be the total number of agents.
 - Let Ψ be the set of all permutations over agents such that every $\psi \in \Psi$ only permutes agents (I, k) assigned to the same information set I .⁷
- The set \widetilde{H} of possible histories is then defined as follows: if $h = (a_1, \dots, a_k) \in H$, then for all $\psi \in \Psi$, $\widetilde{h} = (\psi, a_1, \dots, a_k) \in \widetilde{H}$. We also add the empty sequence ϕ to \widetilde{H} .
- The player assigned to nonterminal history \widetilde{h} is (recall that player “c” is Nature)

$$\widetilde{P}(\widetilde{h}) = \begin{cases} c & \text{if } \widetilde{h} = \phi, \\ c & \text{if } \widetilde{h} = (\psi, h) \text{ and } P(h) = c, \\ (I, \psi(k)) & \text{if } \widetilde{h} = (\psi, h) \text{ and } h \in \widehat{I}_k. \end{cases}$$

- Nature randomizes actions as follows. Given history \widetilde{h} at which Nature operates and action a available at \widetilde{h} ,

$$\widetilde{\rho}(\widetilde{h})(a) = \begin{cases} \frac{1}{|\Psi|} & \text{if } \widetilde{h} = \phi, \\ \rho(h)(a) & \text{if } \widetilde{h} = (\psi, h). \end{cases}$$

- The utility of agent $i = (I, k) \in \widetilde{N}$ at terminal history $\widetilde{h} = (\psi, h)$, is $\widetilde{u}_i(\widetilde{h}) = u_{P(I)}(h)$.
- $\widetilde{\mathcal{I}}$ is the partition with elements $\{\widetilde{h} : \widetilde{P}(\widetilde{h}) = i\}$.

The interpretation is as follows. There is an agent for every frontier of every information set of the original game. If $(I, k) \in \widetilde{N}$, then this agent is assigned to the k -th frontier of information set I . The permutation ψ corresponds to an assignment of the agents to the frontiers within each information set. Nature begins the game by choosing a permutation ψ uniformly at random. This selection fixes the assignment of the agents, however agents only know the information set of the original game at which they operate, but not the frontier assigned to them. Then play resumes as in the original game, with agents making decisions on behalf of the original player, except when Nature is playing in the original game (in which case Nature continues to play the same way in the transformed game). To illustrate, Fig. 4 presents the multiselves agent form of the game of the absentminded driver and the policeman of Section 2. This instance of multiselves agent form is similar to the game presented in Gilboa (1997), and can be interpreted in the same way, the only difference being the

⁷ Formally, $\psi \in \Psi$ if, and only if, for every agent $i = (I, k)$, $\psi(i) = (I', k') \implies I = I'$.




Fig. 4. The multiselves agent form of the absentminded driver and the policeman.

		Agent 2	
		C_2	E_2
Agent 1	C_1	$(1, 1, -1)$	$(0, 0, 0)$
	E_1	$(0, 0, 0)$	$(0, 0, 0)$

Agent 3 plays E

		Agent 2	
		C_2	E_2
Agent 1	C_1	$(0, 0, 0)$	$(\frac{1}{2}, \frac{1}{2}, -\frac{1}{2})$
	E_1	$(\frac{1}{2}, \frac{1}{2}, -\frac{1}{2})$	$(1, 1, 1)$

Agent 3 plays H

Fig. 5. The multiselves agent normal form of the absentminded driver and the policeman.

addition of a player, now represented by Agent 3. It is worth observing that the multiselves agent form has perfect recall since each agent only has one information set, and this information set is never visited twice in the course of the game. In addition, when the original game does not have absentmindedness, Ψ is the trivial subgroup, and thus the transformed game collapses to the standard agent form (with a redundant move by Nature at the beginning of the game).⁸

In the sequel, we use the term “player” to refer to an individual playing the original game and the term “agent” to refer to an individual playing the multiselves agent form (on behalf of some player).

4.2. Partial symmetry

Of course, because the agents who represent the same player operate independently, the multiselves agent game allows for strategy profiles not allowed in the original game.

In this section we establish an equivalence between the strategies of the original game and the strategies of the multiselves agent form that satisfy some symmetry property. To do so, it is convenient to work on games in their normal form. Any extensive game with perfect recall has a normal form representation. We refer to the canonical normal form G of an extensive game Γ as the strategic game with the same set of players in which (i) player i 's action set is composed of the set of pure strategies for player i in game Γ , and (ii) utility functions are induced from the utility function of game Γ .

A *symmetry* of a strategic game G is a permutation over the players of G such that the strategy spaces of permuted players remain the same, and the utility of a player under an action profile is the same as the utility of the permuted player under the permuted action profile.⁹ For example, the permutation $(1\ 2)(3)$, here written in the standard cycle notation, is a symmetry of the strategic game of Fig. 5, while the permutation $(1\ 2\ 3)$ is not. Given a subgroup H of the symmetric group over $\{1, \dots, n\}$, we say that a game G with n players is *H-symmetric* if, for every $\chi \in H$, χ is a symmetry of the game G . In the example of Fig. 5, the only nontrivial subgroup that makes the game *H-symmetric* is the set of two permutations $\{(1)(2)(3), (1\ 2)(3)\}$. An *H-symmetric strategy profile* of an *H-symmetric* game G is a strategy profile σ such that for all $\chi \in H$, $\chi(\sigma) = \sigma$. This notion of symmetry is a major instrument of our existence result, because as we will argue in Proposition 1 below, the existence of a multiselves equilibrium in an extensive game reduces to the existence of

⁸ What is important for our proof to work is that in the multiselves agent form, the partition of information sets does not exhibit absentmindedness, so that the game is one of perfect recall, which allows us to leverage powerful existence techniques inspired by Nash's work. Our choice is conservative: we use the coarsest refinement of the classical agent form that does not have absentmindedness. We find this choice natural, it conforms to the interpretation we make of a player behavior. In particular, the multiselves agent form reduces to the agent form in games without absentmindedness, as mentioned.

⁹ This notion of symmetry is inspired by Nash's seminal work showing the existence of symmetric equilibria (Nash, 1951).

an adequately symmetric equilibrium of the strategic game representing its agent form. This reduction is what renders the problem of existence tractable.

Let the *multiselves agent normal form* be defined as the canonical normal form of the multiselves agent form. Fig. 5 shows the multiselves agent normal form of the game of the absentminded driver and the policeman of Section 2, whose multiselves agent (extensive) form is represented in Fig. 4. Two observations deserve mention. First, the multiselves agent normal form is Ψ -symmetric. Second, there is a bijection between Ψ -symmetric strategy profiles of the multiselves agent normal form and strategy profiles of the original game. This fact follows immediately from that any strategy profile of the multiselves agent form respects Ψ if and only if agents assigned to the same information set of the original game are using the same strategy. Thus, there is a well-defined distribution over actions for each information set, together these distributions defines a strategy profile for the original game. For a strategy profile σ of the original game, we write $\tilde{\sigma}$ the corresponding Ψ -symmetric strategy profile of the multiselves agent form, and inversely, for a Ψ -symmetric strategy profile $\tilde{\sigma}$ of the multiselves agent form, we write σ the corresponding strategy profile of the original game.

The following lemma establishes a bridge between the utility of the multiselves agent (normal) form and that of the original game. Recall that \tilde{u}_i is the utility of agent i in the multiselves agent form.

Lemma 1. *Let σ be a strategy profile for the original game, let μ be a belief consistent with σ , let I be an information set with positive probability of being reached under σ , and let $i = (I, k)$ be an agent of the multiselves agent form assigned to I . Furthermore, let p and p' be distributions over the actions available to play at I , and let ℓ be the player who plays at I . We have*

$$\tilde{u}_i(p, \tilde{\sigma}_{-i}) \geq \tilde{u}_i(p', \tilde{\sigma}_{-i}) \quad \text{if and only if} \quad \text{MSU}_\ell(\sigma; I, p, \mu) \geq \text{MSU}_\ell(\sigma; I, p', \mu).$$

Proof.

$$\begin{aligned} &\tilde{u}_i(p, \tilde{b}_{-i}) \geq \tilde{u}_i(p', \tilde{b}_{-i}) \\ \iff &\frac{1}{D(I)} \sum_{k=1}^{D(I)} u_\ell(\sigma[\hat{I}_k/p]) - \frac{1}{D(I)} \sum_{k=1}^{D(I)} u_\ell(\sigma[\hat{I}_k/p']) \geq 0, \\ \iff &\sum_{k=1}^{D(I)} \sum_{z \in Z} \left(\pi_{\sigma[\hat{I}_k/p]}(z) u_\ell(z) - \pi_{\sigma[\hat{I}_k/p']} (z) u_\ell(z) \right) \geq 0, \\ \iff &\sum_{k=1}^{D(I)} \sum_{h \in \hat{I}_k} \pi_\sigma(h) \sum_{z \in Z} \left(\pi_{\sigma[h/p]}(z | h) u_\ell(z) - \pi_{\sigma[h/p']} (z | h) u_\ell(z) \right) \geq 0, \\ \iff &\frac{1}{\sum_{h \in I} \pi_\sigma(h)} \sum_{h \in I} \pi_\sigma(h) \sum_{z \in Z} \left(\pi_{\sigma[h/p]}(z | h) u_\ell(z) - \pi_{\sigma[h/p']} (z | h) u_\ell(z) \right) \geq 0, \\ \iff &\sum_{h \in I} \mu_I(h) \sum_{z \in Z} \left(\pi_{\sigma[h/p]}(z | h) u_\ell(z) - p_{\sigma[h/p']} (z | h) u_\ell(z) \right) \geq 0, \\ \iff &\sum_{h \in I} \mu_I(h) \left(u_\ell(\sigma[h/p] | h) - u_\ell(\sigma[h/p'] | h) \right) \geq 0, \\ \iff &\text{MSU}_\ell(\sigma; I, p, \mu) \geq \text{MSU}_\ell(\sigma; I, p', \mu). \quad \square \end{aligned}$$

It is not sufficient to simply show the existence of a perfect equilibrium that is also partially symmetric. Instead we must incorporate partial symmetry as part of the notion of perfection. To this end we define the notion of a partially symmetric perturbation. A perturbation η is considered an *H-symmetric perturbation* if for all $\chi \in H$ $\chi(\eta) = \eta$. Note that, in this definition, we again abuse notation and apply permutations to perturbations, analogously to the case of strategy profiles. Similarly, an *H-symmetric perturbed game* is a perturbed game (G, η) where η is an *H-symmetric perturbation*.

Let G be an *H-symmetric normal-form game*. The strategy profile σ is an *H-symmetric perfect equilibrium* if there exists a sequence of *H-symmetric perturbed normal-form games* $(G, \eta_1), (G, \eta_2), \dots$ and a sequence of *H-symmetric Nash equilibria* $\sigma^1, \sigma^2, \dots$ of these perturbed games such that $\sigma^k \rightarrow \sigma$ and $\eta_k \rightarrow 0$.

The next proposition establishes an equivalence between the multiselves perfect equilibria of the original game, and the partially symmetric perfect equilibria of its multiselves agent (normal) form. It implies that proving existence of a multiselves perfect equilibrium of the original game reduces to proving existence of a partially symmetric perfect equilibrium of the multiselves agent normal form.

Proposition 1. *A strategy profile for game Γ is a multiselves perfect equilibrium if and only if it corresponds to a Ψ -symmetric perfect equilibrium of the multiselves agent normal form of Γ .*

Proof. The one-to-one mapping between Ψ -symmetric strategy profiles of the multiselves agent normal form and strategy profiles of the original game also holds for perturbations: there is a bijection between Ψ -symmetric perturbations of the multiselves agent form and perturbations of the original game. As with strategy profiles, for a perturbation of the original game η , we write $\tilde{\eta}$ the corresponding Ψ -symmetric perturbation of the multiselves agent form.

Let σ be a perturbed strategy profile of the original game (for perturbation η) with consistent belief μ , and $\tilde{\sigma}$ the corresponding (perturbed) strategy profile of the multiselves agent form (for perturbation $\tilde{\eta}$).

If the inequality $\tilde{u}_i(\tilde{\sigma}) \geq \tilde{u}_i(p', \tilde{\sigma}_{-i})$ holds for every agent $i = (I, k)$ and every perturbed deviation p' for agent i , then $\tilde{\sigma}$ is a Nash equilibrium of the perturbed multiselves agent form. Moreover, it is a Ψ -symmetric Nash equilibrium, because $\tilde{\sigma}$ is a Ψ -symmetric strategy profile.

Likewise, if the inequality $\text{MSU}_\ell(\sigma; I, \sigma(I), \mu) \geq \text{MSU}_\ell(\sigma; I, p', \mu)$ holds for every player ℓ , every information set I at which ℓ operates, and every perturbed deviation p' , then σ is an MAE of the original perturbed game.

Hence, by Lemma 1, σ is an MAE of the perturbed extensive game (Γ, η) if and only if $\tilde{\sigma}$ is a Nash equilibrium of the perturbed strategic game $(\tilde{\Gamma}, \tilde{\eta})$. It follows that a strategy profile σ is an MPE of game Γ if and only if it corresponds to a Ψ -symmetric perfect equilibrium of the multiselves agent normal form of Γ . \square

4.3. Existence of a partially symmetric equilibrium

We now prove existence of a partially symmetric perfect equilibrium of the multiselves agent normal form. The next lemma, whose proof is relegated to the appendix, asserts existence of a partially symmetric Nash equilibrium in a partially symmetric perturbed strategic game.

Lemma 2. *Let G be an H -symmetric game and η an H -symmetric perturbation of G . There exists an H -symmetric Nash equilibrium of (G, η) .*

It is easily verified that Lemma 2 continues to hold for case $\eta = 0$,¹⁰ which implies existence of partially symmetric equilibria in partially symmetric strategic games, a result which may be of independent interest.

Corollary 1. *Every H -symmetric game has an H -symmetric Nash equilibrium.*

The next proposition asserts existence of a partially symmetric perfect equilibrium in every partially symmetric strategic game.

Proposition 2. *Let G be an H -symmetric strategic game. There exists an H -symmetric perfect equilibrium of G .*

Proof. Let η_1, η_2, \dots be an arbitrary sequence of H -symmetric perturbations that converges to 0. For each perturbation η_k , let σ^k an H -symmetric Nash equilibrium, whose existence is guaranteed by Lemma 2. As the set of H -symmetric strategy profiles is a compact set of strategy profiles, there exists a subsequence of the sequence $\sigma^1, \sigma^2, \dots$ that converges to an H -symmetric strategy profile. This strategy profile satisfies the conditions required of an H -symmetric perfect equilibrium. \square

Theorem 1 then follows from the two key results stated in Propositions 1 and 2.

5. Other equilibrium properties

In this section we present other properties that multiselves solution concepts satisfy. The first property is inclusion, which holds irrespective of whether the game is a game of perfect or imperfect recall.

Proposition 3. *In every game (with or without perfect recall) the following statements hold:*

- (i) *Every multiselves perfect equilibrium is a multiselves sequential equilibrium.*
- (ii) *Every multiselves sequential equilibrium is a multiselves agent equilibrium.*

The proof of Proposition 3 is in Appendix B.2. An analogous result holds for the multiselves Nash equilibrium presented in Appendix A.

One important implication of Proposition 3 and Theorem 1 is equilibrium existence for all multiselves solution concepts.

¹⁰ Strictly speaking, the case $\eta = 0$ is not a special case of Lemma 2 because a well-defined perturbation puts a positive weight on every action at every information set.




Fig. 6. (R_1, R_2) is an agent equilibrium but is not a Nash equilibrium.




Fig. 7. (R, r) is a Nash equilibrium but is not sequential equilibrium.




Fig. 8. (R, r) is a sequential equilibrium but is not a perfect equilibrium.

Corollary 2. All games of imperfect recall have a multiselves sequential equilibrium and an agent equilibrium.

The multiselves equilibrium notions are motivated by the nonexistence of the more traditional equilibria in games of imperfect recall, along with the conceptual barriers, paradoxes and inconsistencies that emerge in those games when one assumes players can behave exactly as if the game was of perfect recall. Clearly, when a game is of imperfect recall, the outcomes of the two classes of solution concepts, multiselves versus traditional, may diverge since existence is not guaranteed for traditional equilibria. However, these notions remain well defined for games of perfect recall. One may ask if the predictions made by multiselves solution concepts remain consistent for games of perfect recall, for which existence is already guaranteed under the traditional definitions of equilibrium and no inconsistency arises. Our next result states that the outcomes of the multiselves solution concepts are exactly the same as the outcomes of their traditional counterparts.

Proposition 4. In every game with perfect recall Γ , the following statements hold:

- (i) A strategy profile is a multiselves agent equilibrium of Γ if and only if it is a Nash equilibrium of the agent form of Γ .
- (ii) A strategy profile is a multiselves perfect equilibrium of Γ if and only if it is a perfect equilibrium of Γ .
- (iii) A strategy profile is a multiselves sequential equilibrium of Γ if and only if it is a sequential equilibrium of Γ .

The proof of Proposition 4 is in Appendix B.3, and, once again, an analogous result holds for the multiselves Nash equilibrium presented in Appendix A.

An immediate consequence of Proposition 4 (and its analog Proposition 7 for the multiselves Nash equilibrium) is that the inclusions of Proposition 3 (and its analog Proposition 6) are strict, observing that the inclusion is strict for the traditional solution concepts: Figs. 6, 7, and 8 show games for which strict inclusion is satisfied.

6. Concluding remarks

We have extended the classical solution concepts for extensive games with perfect recall to games with imperfect recall by integrating the idea of modified multiselves of Piccione and Rubinstein. We have shown that, like the classical equilibrium notions, the multiselves equilibrium notions exist in every game, form a strict hierarchy, and that they collapse to their classical counterparts in games of perfect recall.

To demonstrate these equilibrium properties we have defined novel notions such as the multiselves agent form, phantom strategies, and partial symmetry, which may be of interest beyond the scope of this paper. On a conceptual level, this paper formalizes the idea according to which an individual player can be broken down into basic units, or agents, that are each associated to a node within an information set. Given a state of knowledge, captured by the information set at a decision point, the player is made up of multiple “phantom” independent decision makers, each of whom operates at the level of a single node and has no control beyond the particular instance at which he operates, as in the formulation of Gilboa (1997).

Our results rest on certain assumptions on player behavior. The most prominent assumption is that a player finds herself unable to determine her future actions beyond the immediate choice she confronts. While the assumption is supported by a number of authors, it remains a specific interpretation of games of imperfect recall that may not fit all applications, as Rubinstein (1998) points out in his Chapter 4. Our model also assumes that, given a strategy profile, a player holds subjective probabilities of a node within an information set proportionally to the probability at which the node is visited under the said profile. While we inherit this assumption from previous models—in particular Piccione and Rubinstein (1997)—it remains substantive and other definitions of subjective probability are sensible. We do not know if our results hold for other definitions of subjective probability, it appears to us that such a change entails a quite different way of proof, and we leave this point for future research. Similarly, we have made a specific use of phantom strategies, namely we have used phantom strategies for nodes located at the frontiers of an information set. This choice appeared to us as a natural one, yet as we observed in Section 4, our results do not rely on this assumption, since they also hold with other subpartitions of information sets, as long as those do not exhibit absentmindedness.

Appendix A. On the multiselves Nash equilibrium

In this Appendix, we propose a definition of multiselves Nash equilibrium, and we show that this equilibrium notion satisfies the same properties as the other three equilibrium notions introduced in the main text, namely, that a multiselves Nash equilibrium is guaranteed to exist in all games, and that it satisfies the inclusion and reduction properties analogously to Propositions 3 and 4.

We start by our definition of multiselves Nash equilibrium. This definition uses the notion of equivalent strategy profiles defined by Battigalli (1997). Two strategy profiles σ and σ' are equivalent, written $\sigma \sim \sigma'$, if they induce the same probability for every history. Beyond strategy profiles, the definition extends directly to individual player strategies: given a strategy profile σ , two strategies σ'_i and σ''_i of player i are said to be equivalent with respect to σ , and written $\sigma'_i \sim_\sigma \sigma''_i$, if $(\sigma'_i, \sigma_{-i}) \sim (\sigma''_i, \sigma_{-i})$.

Definition 5. A strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is a *multiselves Nash equilibrium* (MNE) if it is a multiselves agent equilibrium and for every player i there exists a strategy $\sigma'_i \sim \sigma_i$, a sequence of completely mixed strategies $\sigma_i^1, \sigma_i^2, \dots$ for player i that converges to σ'_i , and a sequence of positive reals $\varepsilon_1, \varepsilon_2, \dots$ that converges to 0 such that for every k , for every belief μ consistent with strategy profile $\sigma^k = (\sigma_i^k, \sigma_{-i})$, for every information set I assigned to player i with positive probability in σ^k , and for every distribution p over actions available at I , $MSU_i(\sigma^k; I, p, \mu) \leq SU_i(\sigma^k; I, \mu) + \varepsilon_k$.

The definition bears some resemblance with our definition of a sequential equilibrium. Observe however an important difference: the sequence of mixed strategies is composed of individual player strategies as opposed to strategy profiles. Although this definition of multiselves Nash equilibrium may, at first sight, appear distant from the traditional notion of Nash equilibrium—and significantly more complex—we show below that the two notions of equilibrium agree on games of perfect recall.

At high level, what causes this complexity is the fact that the multiselves approach, which commands players to evaluate and decide on immediate actions at all points at which they operate, is not naturally compatible with the idea of Nash equilibrium, which, on the contrary, has a player decide and evaluate her entire plan of action from an ex ante perspective, before even entering the game. As we show, the two ideas can be still be made compatible, yet it remains that our notion of multiselves Nash equilibrium is not as natural as the other three solution concepts. It is included here for completeness. Although there may exist other ways to define a multiselves Nash equilibrium, for the reason just mentioned, we believe that there is no natural definition.

We now state our existence result.

Proposition 5. *All games of imperfect recall have a multiselves Nash equilibrium.*

Proposition 5 follows from Theorem 1, from the fact that every multiselves perfect equilibrium is also a multiselves sequential equilibrium as shown in Proposition 3, and from the fact that every multiselves sequential equilibrium is also a multiselves Nash equilibrium as shown in the following proposition.

Proposition 6. *In every game (with or without perfect recall) the following statements hold:*

- (i) *Every multiselves sequential equilibrium is a multiselves Nash equilibrium.*
- (ii) *Every multiselves Nash equilibrium is a multiselves agent equilibrium.*

Proof. Statement (ii) is true by definition. Let us prove statement (i). First, observe that an MSE is also an MAE and thus the fact that an MNE is defined to be an MAE plays no part in the remaining proof—it follows directly from the definitions.

Now, Let σ be an MSE with corresponding sequence of strategy profiles $\sigma^1, \sigma^2, \dots$ and corresponding sequence of positive reals $\varepsilon_1, \varepsilon_2, \dots$. Since strategy equivalence is an equivalence relation, $\sigma_i \cong_{\sigma} \sigma_i$ for all i . For player i , consider the sequence $\sigma_i^1, \sigma_i^2, \dots$ (the sequence of player i strategies drawn from the original sequence of strategy profiles) and the sequence $\varepsilon_1, \varepsilon_2, \dots$. The two sequences just defined satisfy the conditions that define an MNE. \square

The following proposition guarantees that the notion of multiselves Nash equilibrium agrees with the standard notion of Nash equilibrium for games of perfect recall.

Proposition 7. *If Γ is a game with perfect recall, then a strategy profile is a multiselves Nash equilibrium of Γ if and only if it is a Nash equilibrium of Γ .*

Proof. It follows from the definition of an MNE that a strategy profile σ is an MNE if and only if, for every player i , there exists an MSE of player i 's induced decision problem that is equivalent to player i 's strategy in σ (player i 's induced decision problem is the game in which all other players' actions are regarded as fixed acts of Nature). By Proposition 4 this observation leads to the following characterization of an MNE: a strategy profile σ is an MNE if, for every player i , there exists an SE of player i 's induced decision problem, σ'_i , such that $\sigma_i \cong_{\sigma} \sigma'_i$. Additionally, a strategy profile is a Nash equilibrium if and only if every player is playing an optimum in her respective decision problem. Thus, to prove the proposition, it is enough to show that a strategy is an optimum of an extensive decision problem if and only if it is equivalent to a sequential equilibrium.

If part: An SE involves choosing optimal strategies at all initial information sets (information sets that only Nature can precede), and thus for games of perfect recall a player following an SE strategy also follows an optimal strategy ex ante. Furthermore, because two equivalent strategies induce the same distribution over histories, they must have the same payoff and if one is optimal, the other must be as well. Hence, if a strategy is equivalent to a sequential equilibrium, it is optimal.

Only If part: It follows from Propositions 3.5 and 4.3 of Battigalli (1997) that for every optimal strategy σ of a decision problem, there exists an MSE σ' such that $\sigma \cong \sigma'$.¹¹ Because the notion of MSE reduces to the notion of SE in games of perfect recall, for every optimal strategy there exists an equivalent SE. \square

Appendix B. Proofs omitted from the main text

B.1. Proof of Lemma 2

Let G be an H -symmetric strategic game and η be an H -symmetric perturbation of G .

We begin by defining a specific set of strategies for player i , $\{p_{ia} : a \in A(i)\}$, as

$$p_{ia} = \left(1 - \sum_{a' \in A(i)} \eta(a') \right) \cdot \mathbb{1}_a + \sum_{a' \in A(i)} \eta(a') \cdot \mathbb{1}_{a'}.$$

In words, p_{ia} is the strategy for player i with the largest possible weight given to action a allowed in the perturbed game. It is a vertex of the strategy simplex for player i , and the convex hull of the set $\{p_{ia} : a \in A(i)\}$ is the set of possible strategies of player i allowed in the perturbed game.

Next, let σ be a strategy profile of the perturbed game, i be a player, and a be an action for that player. We define the continuous functions ϕ_{ia} of σ by

¹¹ Battigalli actually proves this fact for what he calls a modified multiselves sequential equilibrium, which is equivalent to a single player version of an MSE in this paper.

$$\phi_{ia}(\sigma) = \max(0, u_i(p_{ia}, \sigma_{-i}) - u_i(\sigma)).$$

We also define continuous functions T_i on the individual components σ_i of σ as

$$T_i(\sigma_i) = \frac{\sigma_i + \sum_{a \in A(i)} \phi_{ia}(\sigma) \cdot p_{ia}}{1 + \sum_{a \in A(i)} \phi_{ia}(\sigma)}.$$

Finally, we let $T(\sigma) = (T_1(\sigma_1), \dots, T_n(\sigma_n))$.

The mapping T is a variation of Nash's function T (Nash, 1951). One important difference in our version of T is that T is a continuous mapping from the strategy space of the game perturbed by η to itself, instead of the full strategy space (note that $T(\sigma)$ is a convex combination of two perturbed game strategy profiles). However, the fact that a strategy profile σ is a fixed point of T if and only if it is a Nash equilibrium of the perturbed game continues to hold. This fact owes to the other fact that a strictly better response for player i exists if and only if one of the vertices of player i 's strategy simplex is a strictly better response.

As Nash noted, convex combinations of strategies that respect a particular symmetry χ also respect χ , thus the set of H -symmetric strategies forms a compact convex set, as the intersection of compact convex sets. Furthermore, for every $\chi \in H$, we have $\phi_{\chi(i)a}(\chi(\sigma)) = \phi_{ia}(\sigma)$. So $\chi(T(\chi(\sigma))) = T(\sigma)$ and thus T applied to an H -symmetric strategy produces another H -symmetric strategy. Besides, if σ is a strategy profile of the perturbed game, $\chi(T(\sigma))$ is a valid perturbed game strategy profile for all $\chi \in H$, because η is an H -symmetric perturbation. Therefore, T is a continuous mapping from the set of H -symmetric strategies of the perturbed game to itself and by Brouwer's fixed point theorem, T has a fixed point in H -symmetric strategies. We conclude that there exists an H -symmetric Nash equilibrium of the game (G, η) .

B.2. Proof of Proposition 3

The fact that a multiselves sequential equilibrium is also a multiselves agent equilibrium follows from the definitions. Let us show that a multiselves perfect equilibrium is also a multiselves sequential equilibrium.

Let σ be an MPE with corresponding sequence of perturbations η_1, η_2, \dots and corresponding sequence of perturbed MAE $\sigma^1, \sigma^2, \dots$. Given a perturbation η , an information set I (assigned to player i), and a distribution $p \in \Delta(A(I))$, consider the distribution-valued mapping f_η defined as

$$f_\eta(I, p) = \eta(I) + \left(1 - \sum_{a \in A(I)} \eta(I)(a)\right) \cdot p.$$

This mapping has the effect of "shifting" the distribution p in the direction specified by $\eta(I)$. Importantly, if σ is a perturbed strategy profile for perturbation η , then for all $h \in I$, $\sigma[h/f_\eta(I, p)]$ is a valid perturbed phantom strategy profile. Moreover, $f_\eta(I, p)$ is continuous with respect to η , and if $\eta = 0$ it maps to the initial distribution. So, as $\eta_k \rightarrow 0$, we can choose some positive integer K such that for all $k \geq K$, all strategy profiles σ , all information sets I assigned to player i , and all $p \in \Delta(A(I))$,

$$\left|u_i(\sigma[h/f_{\eta_k}(I, p)] | h) - u_i(\sigma[h/p] | h)\right| < \varepsilon.$$

Now let $\varepsilon_1, \varepsilon_2, \dots$ be a sequence of positive reals that converges to 0. Let ℓ be an arbitrary positive integer. By the above argument and the definition of multiselves subjective utility, we can choose a positive integer K_ℓ strictly increasing in ℓ and such that for all $k \geq K_\ell$, all beliefs μ consistent with σ^k , all information sets I assigned to player i , and all $p \in \Delta(A(I))$,

$$\left|MSU_i(\sigma^k; I, f_{\eta_k}(I, p), \mu) - MSU_i(\sigma^k; I, p, \mu)\right| < \varepsilon.$$

Additionally, as σ^k is an MAE of (Γ, η_k) and $\sigma^k[h/f_{\eta_k}(I, p)]$ is a valid perturbed phantom strategy profile for all $h \in I$, we have

$$MSU_i(\sigma^k; I, f_{\eta_k}(I, p), \mu) \leq SU_i(\sigma^k; I, \mu).$$

Combining the last two inequalities yields

$$MSU_i(\sigma^k; I, p, \mu) \leq SU_i(\sigma^k; I, \mu) + \varepsilon.$$

As K_ℓ is strictly increasing in ℓ , the sequence of completely mixed strategies $\sigma_{K_1}, \sigma_{K_2}, \dots$ converges to σ . This sequence of mixed strategies, along with the sequence of positive reals $\varepsilon_1, \varepsilon_2, \dots$, satisfy the conditions required in an MSE, and so σ is an MSE.

B.3. Proof of Proposition 4

This proof makes use of the following lemma.

Lemma 3. *A strategy profile is a multiselves agent equilibrium of game Γ if and only if it corresponds to a Ψ -symmetric Nash equilibrium of the multiselves agent normal form of Γ .*

Proof. The argument is similar to the proof of Proposition 1. In Section 4.2 we established a one-to-one mapping between Ψ -symmetric strategy profiles of the multiselves agent form of Γ and strategy profiles of Γ . Let σ be a strategy profile of Γ with consistent belief μ and $\tilde{\sigma}$ be the corresponding strategy profile of the multiselves agent form of Γ .

If σ is a multiselves agent equilibrium of game Γ , then the inequality $MSU_\ell(\sigma; I, \sigma(I), \mu) \geq MSU_\ell(\sigma; I, p', \mu)$ holds for every player ℓ , every information set I at which ℓ operates and that is reached with positive probability, and every deviation p' . By Lemma 1, the inequality $\tilde{u}_i(\tilde{\sigma}) \geq \tilde{u}_i(p', \tilde{\sigma}_{-i})$ also holds, for every agent $i = (I, k)$ for which I is reached with positive probability under σ , and every deviation p' for agent i . Note that if agent i plays at information set I , but I is reached with zero probability under σ , then agent i 's payoffs are the same: for all pairs of distribution p, p' over the possible actions of agent i , $\tilde{u}_i(p, \tilde{\sigma}_{-i}) = \tilde{u}_i(p', \tilde{\sigma}_{-i})$. Hence, by definition, $\tilde{\sigma}$ is a Ψ -symmetric Nash equilibrium of the multiselves agent normal form of Γ .

Conversely, if $\tilde{\sigma}$ is a Ψ -symmetric Nash equilibrium of the multiselves agent normal form of Γ , then the inequality $\tilde{u}_i(\tilde{\sigma}) \geq \tilde{u}_i(p', \tilde{\sigma}_{-i})$ holds for every agent $i = (I, k)$ and every deviation p' for agent i , and so by Lemma 1, the inequality $MSU_\ell(\sigma; I, \sigma(I), \mu) \geq MSU_\ell(\sigma; I, p', \mu)$ also holds for every player ℓ , every information set I at which ℓ operates and that is reached with positive probability under σ , and every deviation p' . Hence, by definition, σ is a multiselves agent equilibrium of game Γ . \square

Note that, in games without absentmindedness (which include games of perfect recall), Ψ is the trivial subgroup, and the multiselves agent form is simply the standard agent form of the game. Then Lemma 3 immediately implies that in games without absentmindedness, a strategy profile is a multiselves agent equilibrium if and only if it is a Nash equilibrium of the agent form—and hence an agent equilibrium. This fact establishes statement (i) of Proposition 4.

By a similar argument, Proposition 1 implies that in games without absentmindedness, a strategy profile is a multiselves perfect equilibrium if and only if it is a perfect equilibrium of the agent form. By Theorem 4 of Selten (1975), for games of perfect recall, a strategy profile is a perfect equilibrium of the agent form if and only if it is a perfect equilibrium.¹² This fact establishes statement (ii) of Proposition 4.

The case of sequential equilibrium is more complicated, and requires to introduce several intermediary results, and a definition. Given a strategy profile σ , a belief μ consistent with σ , a player i , an information set I at which player i operates, and a deviation $p \in \Delta(A(I))$, the subjective utility of this deviation is given by

$$SU_i(\sigma; I, p, \mu) = SU_i((\sigma_i[I/p], \sigma_{-i}); I, \mu),$$

where $\sigma_i[I/p]$ is the uniform-deviation strategy as defined in Section 3.2.

Lemma 4. *Let σ be a strategy profile of a game without absentmindedness, and μ be a belief consistent with σ . For every player i who operates at information set I and for every distribution p over the actions available at I ,*

$$SU_i(\sigma; I, p, \mu) = MSU_i(\sigma; I, p, \mu).$$

Proof. The proof follows from the observation that, in games without absentmindedness, for any $h \in I$, $u_i(\sigma[I/p] | h) = u_i(\sigma[h/p] | h)$. \square

Lemma 5. *Let σ be a strategy profile of a game with perfect recall Γ , and μ be a belief consistent with σ . Let I be an information set at which player i operates, and \mathcal{J} be the set of all information sets that come after I and are still controlled by player i .¹³ If the positive reals $\varepsilon_J, J \in \mathcal{J}$ are such that for all $J \in \mathcal{J}$ and all strategy profiles σ' ,*

$$SU_i((\sigma'_i, \sigma_{-i}); J, \mu) \leq SU_i(\sigma; J, \mu) + \varepsilon_J,$$

and if the positive real ε_I is such that for all $a \in A(I)$,

$$SU_i(\sigma; I, a, \mu) \leq SU_i(\sigma; I, \mu) + \varepsilon_I,$$

¹² This result was actually proved for the agent normal form of a game, but as the Nash equilibria of the agent form and the agent normal form coincide, even for perturbed variants, it is equivalent.

¹³ This set of information sets is well defined because the game does not have absentmindedness.

then for all strategy profiles σ' ,

$$SU_i((\sigma'_i, \sigma_{-i}); I, \mu) \leq SU_i(\sigma; I, \mu) + \varepsilon_I + \max_{Y \in J} \varepsilon_Y,$$

where by convention $\max_{Y \in J} \varepsilon_Y = 0$ if $J = \emptyset$.

Proof. Assume the two conditions of Lemma 5 are satisfied. Fix an arbitrary action $a \in A(I)$. Let $\widehat{\mathcal{J}}$ be the set of information sets that are reached after action a is taken at I . (This set is well defined because Γ has perfect recall.) Slightly abusing notation, let $\pi_\sigma(J | I)$ denote the probability of reaching information set $J \in \widehat{\mathcal{J}}$ from I under σ . Fix an arbitrary strategy profile σ' . By the first condition of Lemma 5,

$$SU_i((\sigma'_i, \sigma^k_{-i}); J, \mu) \leq SU_i(\sigma; J, \mu) + \varepsilon_J,$$

so the most that changing strategy at information set J can contribute to increasing the payoff is $\pi_\sigma(J | I)\varepsilon_J$. Therefore,

$$\begin{aligned} SU_i((\sigma'_i[I/a], \sigma_{-i}); I, \mu) &\leq SU_i((\sigma_i[I/a], \sigma_{-i}); I, \mu) + \sum_{J \in \widehat{\mathcal{J}}} \pi_\sigma(J | I)\varepsilon_J, \\ &\leq SU_i((\sigma_i[I/a], \sigma_{-i}); I, \mu) + \max_{J \in \mathcal{J}} \varepsilon_J, \\ &\leq SU_i(\sigma; I, \mu) + \varepsilon_I + \max_{J \in \mathcal{J}} \varepsilon_J, \end{aligned}$$

where the second inequality uses the observation that, as the game has perfect recall, $\sum_{J \in \widehat{\mathcal{J}}} \pi_\sigma(J | I) \leq 1$, and the third inequality is obtained applying the second condition of Lemma 5.

For a game with perfect recall, there is always a pure-strategy best response at every information set. As a result, it is sufficient to consider deviations of the form $\sigma'_i[I/a]$ instead of the more general deviations σ'_i . Together with the inequality just obtained, we conclude that

$$SU_i((\sigma'_i, \sigma_{-i}); I, \mu) \leq SU_i(\sigma; I, \mu) + \varepsilon_I + \max_{J \in \mathcal{J}} \varepsilon_J. \quad \square$$

Lemma 6. *If Γ is a game with perfect recall, then σ is a sequential equilibrium of Γ if and only if it is a sequential equilibrium of the agent form of Γ .*

Proof. Of course, any sequential equilibrium is a sequential equilibrium of the agent game because it expands the set of strategies under consideration at every information set. Let us prove the other direction.

Let $\varepsilon > 0$ and let σ be a sequential equilibrium of the agent form with associated sequences of strategies $\sigma^1, \sigma^2, \dots$ and positive reals $\varepsilon_1, \varepsilon_2, \dots$, as in the definition of Section 3.1. Let μ be a belief consistent with σ . Finally, let m be the largest possible number of decisions that can be made by any one player in the course of play (as the game has perfect recall, the number of decisions reduces to the number of information sets the player passes through in a given instance of play). As $\varepsilon_k \rightarrow 0$, there exists an index k such that $\varepsilon_k \leq \varepsilon/m$.

Fix a player i and any information set I at which that player operates. Let d be the maximum possible number of moves left for player i starting from I , including the decision at I . An inductive application of Lemma 5 yields that for all strategy profiles σ' ,

$$\begin{aligned} SU_i((\sigma'_i, \sigma^k_{-i}); I, \mu) &\leq SU_i(\sigma^k; I, \mu) + d\varepsilon/m, \\ &\leq SU_i(\sigma^k; I, \mu) + \varepsilon, \end{aligned}$$

since $d \leq m$.

If a sequence of such positive reals ε is chosen that converges to 0, the resulting sequence of completely mixed strategies σ^k converges to σ . Thus σ is a sequential equilibrium. \square

We can now conclude the proof of statement (iii) of Proposition 4. First, note that if, in the definition of multiselves sequential equilibrium, we replace the term MSU with the comparable term SU for uniform-deviation strategies defined for Lemma 4, then the resulting equilibrium concept is the same as the notion of sequential equilibrium of the agent form. This observation, combined with Lemma 4, implies that the multiselves sequential equilibria are identical to the sequential equilibria of the agent form in games without absentmindedness.¹⁴ Lemma 6 then implies that the multiselves sequential equilibria are the same as the sequential equilibria in games of perfect recall.

¹⁴ This fact was first observed without proof by Battigalli for single agent games (Battigalli, 1997).

References

- Abreu, D., Rubinstein, A., 1988. The structure of Nash equilibrium in repeated games with finite automata. *Econometrica* 56 (6), 1259–1281.
- Aumann, R.J., Hart, S., Perry, M., 1997. The absentminded driver. *Games Econ. Behav.* 20 (1), 102–116.
- Başar, T., Jan Olster, G., 1999. *Dynamic Noncooperative Game Theory*, second ed.. Society for Industrial and Applied Mathematics.
- Battigalli, P., 1997. Dynamic consistency and imperfect recall. *Games Econ. Behav.* 20 (1), 31–50.
- Gilboa, I., 1997. A comment on the absentminded driver paradox. *Games Econ. Behav.* 20 (1), 25–30.
- Grove, A.J., Halpern, J.Y., 1997. On the expected value of games with absentmindedness. *Games Econ. Behav.* 20 (1), 51–65.
- Halpern, J.Y., 1997. On ambiguities in the interpretation of game trees. *Games Econ. Behav.* 20 (1), 66–96.
- Halpern, J.Y., Pass, R., 2016. Sequential equilibrium in games of imperfect recall. In: *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning*, pp. 278–287.
- Hillas, J., Kvasov, D., 2017. Backward Induction in Games Without Perfect Recall. Working paper.
- Isbell, J., 1957. Finitary games. In: *Contributions to the Theory of Games III*. Princeton University Press, pp. 79–96.
- Jehiel, P., 1995. Limited horizon forecast in repeated alternate games. *J. Econ. Theory* 67 (2), 497–519.
- Jehiel, P., 2001. Limited foresight may force cooperation. *Rev. Econ. Stud.* 68 (2), 369–391.
- Johansson, M., Waugh, K., Bowling, M., Zinkevich, M., 2011. Accelerating best response calculation in large extensive games. In: *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*. IJCAI-11.
- Kline, J.J., 2005. Imperfect recall and the relationships between solution concepts in extensive games. *Econ. Theory* 25 (3), 703–710.
- Kreps, D.M., Wilson, R., 1982. Sequential equilibria. *Econometrica* 50 (4), 863–894.
- Kuhn, H., 1953. Extensive games and the problem of information. In: *Dresher, D., Tucker, A.W., Wolfe, P. (Eds.), Contributions to the Theory of Games II*. Princeton University Press, pp. 193–216.
- Lipman, B.L., 1997. More absentmindedness. *Games Econ. Behav.* 20 (1), 97–101.
- Monderer, D., Tennenholtz, M., 1999. Distributed games. *Games Econ. Behav.* 28 (1), 55–72.
- Nash, J., 1951. Noncooperative games. *Ann. Math.* 54 (2), 286–295.
- Neyman, A., 1985. Bounded complexity justifies cooperation in the finitely repeated prisoners' dilemma. *Econ. Lett.* 19 (3), 227–229.
- Osborne, M.J., Rubinstein, A., 1994. *A Course in Game Theory*. MIT Press.
- Papadimitriou, C.H., Yannakakis, M., 1994. On complexity as bounded rationality. In: *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing*, pp. 726–733.
- Piccione, M., Rubinstein, A., 1997. On the interpretation of decision problems with imperfect recall. *Games Econ. Behav.* 20 (1), 3–24.
- Rubinstein, A., 1986. Finite automata play the repeated prisoner's dilemma. *J. Econ. Theory* 39 (1), 83–96.
- Rubinstein, A., 1998. *Modeling Bounded Rationality*. MIT Press.
- Samuelson, L., 2001. Analogies, adaptation, and anomalies. *J. Econ. Theory* 97 (2), 320–366.
- Segal, U., 2000. Don't fool yourself to believe you won't fool yourself again. *Econ. Lett.* 67 (1), 1–3.
- Selten, R., 1975. Reexamination of the perfectness concept for equilibrium points in extensive games. *Int. J. Game Theory* 4 (1), 25–55.
- Strotz, R.H., 1955. Myopia and inconsistency in dynamic utility maximization. *Rev. Econ. Stud.* 23 (3), 165–180.
- Wichardt, P.C., 2008. Existence of Nash equilibria in finite extensive form games with imperfect recall: a counterexample. *Games Econ. Behav.* 63 (1), 366–369.
- Wichardt, P.C., 2010. Modelling equilibrium play as governed by analogy and limited foresight. *Games Econ. Behav.* 70 (2), 472–487.