

Imperfect recall and time inconsistencies: an experimental test of the absentminded driver “paradox”

M. Vittoria Levati · Matthias Uhl · Ro'i Zultan

Accepted: 26 February 2013 / Published online: 23 March 2013
© Springer-Verlag Berlin Heidelberg 2013

Abstract Absentmindedness is a special case of imperfect recall, in which a single history includes more than one decision node in an information set. Put differently, players, after making a decision, sometimes face it again without recalling having ‘been there before’. Piccione and Rubinstein (Game Econ Behav 20(1):3–24, 1997b) have argued that absentmindedness may lead to time inconsistencies. Specifically, in certain cases, a player’s optimal strategy as calculated when called to choose an action (the *action stage*) deviates from the optimal strategy as calculated in a preceding *planning stage*, although preferences remain constant and no new information is revealed between the two stages. An alternative approach assumes that the player maximizes expected payoff in the action stage while considering his actions at other decision nodes to be immutable. With this approach, no time inconsistencies arise. The present paper explores this issue from a behavioral point of view. We elicit participants’ strategies in an experimental game of absentmindedness, separately for a planning stage and an action stage. We find systematic and robust time inconsistencies under four variations of the experiment and using ten different parameterizations of the game.

Electronic supplementary material The online version of this article (doi:[10.1007/s00182-013-0373-y](https://doi.org/10.1007/s00182-013-0373-y)) contains supplementary material, which is available to authorized users.


M. V. Levati
Department of Economics, University of Verona, Via dell’Artigliere 19, 37129 Verona, Italy

M. V. Levati
Max Planck Institute of Economics, Kahlaische Straße 10, 07745 Jena, Germany

M. Uhl
TUM School of Education, Technical University of Munich, Lothstraße 17, 80335 Munich, Germany

R. Zultan (✉)
Department of Economics, Ben-Gurion University of the Negev, P.O.B. 653, Beer-Sheva 84105, Israel
e-mail: zultan@bgu.ac.il

Fig. 1 The absentminded driver problem



We conclude that real decisions under absentmindedness without commitment are susceptible to time inconsistencies.

Keywords Imperfect recall · Absentmindedness · Dynamic inconsistency · Experiment

1 Introduction

Dynamic consistency is a compelling fundamental tenet of rational behavior: once a decision maker makes a plan, he should carry it out as long as there is no relevant change in the decision environment. Notwithstanding its normative appeal, the principle of dynamic consistency has been systematically invalidated by empirical evidence, therefore calling for a revision of the normative theories, as in the case of decision-making under risk (e.g. [Kahneman and Tversky 1979](#)) or ambiguity (e.g. [Gilboa and Schmeidler 1989](#); [Epstein and Schmeidler 2003](#)). Conversely, [Piccione and Rubinstein \(1997b, henceforth PR\)](#) have drawn attention to a particular case of dynamic inconsistency that arises exactly from standard rational decision theory. PR considered a specific type of imperfect recall, which they termed “absentmindedness”, where a single history passes through two decision nodes in an agent’s information set. They showed that considerations that would bear no impact on the analysis of decision problems with perfect recall become crucial to the analysis when absentmindedness (and, to a lesser extent, other forms of imperfect recall) is involved.

Absentmindedness and the paradoxical results associated with it are most usefully illustrated by the problem of the “absentminded driver”, a simplified version of which is presented in Fig. 1’s game tree.¹

The absentminded driver starts his journey at intersection X (the first decision node in the information set), where he can either “exit” (for a payoff of a) or “continue” to intersection Y (the second decision node in the information set), where he faces an identical choice. If at Y he exits, he gets a payoff of b ; if he continues beyond Y , he earns c . The driver suffers from absentmindedness in the sense that he is unable to distinguish between nodes X and Y , both of them being in the same information set.

¹ We speak of game (rather than decision) tree to stay in the framework of game theory.

By using pure strategies, the decision maker can only obtain either payoff a or c . Thus, by using mixed strategies, i.e., randomizing over pure strategies, the decision maker can obtain any convex combination of a and c . It follows that if $b > a, c$, then the highest payoff b can never be reached through either pure or mixed strategies. Nonetheless, a higher expected payoff can be obtained through *behavioral strategies*, by which the decision maker randomizes over his actions independently at X and at Y (Kuhn 1953; Isbell 1957).²

PR demonstrated that, if the highest payoff is at the second exit, then an agent's plan before he starts his journey (the *planning stage*) is inconsistent with his beliefs once he reaches a decision node (the *action stage*) as long as he assigns some positive probability to being at Y . In other words, the decision maker is tempted to change his initial plan when the time comes to execute it. This observation was termed by PR "the absentminded driver paradox".

PR suggested that the problem can be alternatively analyzed by applying a "modified multiself" approach, by which the decision maker is taken as be comprised of two "selves", each independently making a decision for one decision node in the information set. In this interpretation, when the decision maker is called to make a choice, he takes the decision (to be) made by his "twin self" as immutable. In this interpretation of the problem, the symmetric equilibrium, in which the behavioral strategy taken at the information set is a best response to itself, coincides with the planning-optimal strategy so that the paradox vanishes. Thus, the analysis of the problem depends crucially on whether the decision maker can commit to a strategy and whether he is assumed to take his actions in future occurrences of the same information set as under his control.

Whereas PR presented both ways of reasoning about the problem as potentially valid, Aumann et al. (1997a, henceforth AHP) took a more restrictive normative stand. They argued that optimization at the information set *must* be carried out with respect to the strategy at the current decision node while considering the rest of the play as fixed, as in PR's multiself approach. AHP showed that, in this case, the planning-optimal decision is necessarily action-optimal, thus resolving the paradox.

While there exists an extensive normative debate on the subject of absentmindedness,³ little is known about actual behavior in such situations. In fact, absentmindedness is prevalent in real-world situations. Two of the authors' mothers, for instance, take prescriptions to moderate blood pressure on a daily basis. They are often faced with the decision about whether to take a pill or not as they cannot remember if they took the pill that day. This problem is exacerbated if many similar decisions are to be made over time; for example, if there are several medications, each to be administered according to a different schedule. Organizations are also susceptible to

² To illustrate the difference between mixed and behavioral strategies in our game, compare the mixed and the behavioral strategy that both specify a probability of 0.5 to continue. The former implies that the same realized action is chosen at both X and Y , resulting in an expected payoff of $0.5a + 0.5c$. With the latter, the actions chosen at X and Y are determined independently, resulting in an expected payoff of $0.5a + 0.25b + 0.25c$.

³ Theoretical discussions of the paradoxes arising under absentmindedness can be found in Battigalli (1997), Gilboa (1997), Grove and Halpern (1997), Halpern (1997), Aumann et al. (1997b), and Lipman (1997), which are summarized and countered in Piccione and Rubinstein (1997a), Binmore (1996), Kline (2005) and Board (2003). The question of beliefs under absentmindedness spawned a prolific literature in philosophy starting with Elga (2000).

absentmindedness, for example when devising a decentralized policy that will be carried out by different employees who do not know previous decisions of each other, or even whether such decisions were made (Binmore 1996; see also Isbell 1957).

The prevalence of situations involving absentmindedness calls for a better understanding of how the theoretical debate reflects on actual behavior. This paper is aimed at addressing the question of dynamic inconsistency under absentmindedness from an empirical perspective. In particular, we aim to test whether the time inconsistencies that emerge from the theoretical analysis inform actual behavior when people are placed in a situation corresponding to the theoretical problem. That is, we compare behavior in a planning stage and in an action stage of (a modified version of) the absentminded driver problem.⁴

The existing experimental literature contains two studies of the absentminded driver problem. In the first study, Huck and Müller (2002) constructed a situation of absentmindedness by matching pairs of decision makers who, together, constitute the absentminded driver. They essentially implemented the multiself approach as they put different individuals in different decision nodes. However, Huck and Müller (2002) did not elicit planning-stage strategies, but focused on comparing behavior in repeated interactions with and without rematching.

More similar to our experiment is the innovative study by Deck and Sarangi (2009), who created a situation of absentmindedness through information overload to find that action-stage behavior deviates from the theoretically optimal strategy in the direction predicted by the paradox.⁵ To draw conclusions about time inconsistencies from this result, however, one must assume that behavior in the planning stage quantitatively matches the optimal strategy, which is ultimately an empirical question.⁶ Deck and Sarangi (2009) did ask for planning-stage strategies for a small set of parameters and observed a significant difference between the elicited action-stage and planning-stage strategies. However, the participants in their experiment were restricted to choosing between two actions (either to exit or to continue). Hence, in the action stage they were able to implement behavioral strategies by randomizing over the two actions every time they were required to make a decision. Conversely, in the planning stage, where participants chose a single action, which was then applied by the experimenters to both nodes, only mixed strategies were implementable. Since the theoretical predictions differ for behavioral and mixed strategies even within the planning stage (Isbell 1957), this procedure does not provide an empirical test of the theoretical time inconsistencies.⁷ In contrast, participants in our experiment can implement


⁴ Methodologically, the experiment can be taken as a test of the game-theoretical predictions. As such, our goal was to construct a laboratory situation which is in the *base domain* of the theory. Consequently, the universal nature of game theory places the experiment firmly within the *test domain* of the theory (Bardsley et al. 2010).

⁵ We will discuss the details of their design in Sect. 3.

⁶ For example, presenting the decision as one of whether to exit to the *right* may create a bias to exit (cf. Wilson and Nisbett 1978).

⁷ For example, with payoffs $(a, b, c) = (1, 5, 2)$, the optimal mixed strategy is to continue with probability 1 for a payoff of 2, whereas the optimal behavioral strategy is to continue with probability $\frac{2}{3}$ for a payoff of $2\frac{1}{3}$. Note that implementing the planning-optimal behavioral strategy as a mixed strategy yields a sub-optimal payoff of $1\frac{2}{3}$.

Fig. 2 The modified absentminded driver problem



behavioral strategies in both stages through the use of an identical exogenous randomizing mechanism. This allows us to cleanly compare behavior in the two stages.

To sum, our contribution is to provide an empirical comparison of behavior in a planning stage and an action stage in a situation of naturally induced absentmindedness, where behavioral strategies are implemented through an exogenous randomizing device. The hypothesis under consideration is that participants will exit more often in the action stage than in the planning stage, as suggested by PR’s theoretical analysis. The data strongly support this hypothesis under four variations of the experiment and using ten different parameterizations of the game. We conclude that real decisions under absentmindedness without commitment are susceptible to time inconsistencies.

The remainder of the paper is organized as follows. The next section provides a theoretical analysis of our modified experimental game. Section 3 details the experimental design. Section 4 discusses our experimental results, and Sect. 5 concludes.

2 Theoretical analysis of the experimental game

In the original formulation of the game, the decision maker passes through the information set either once or twice, depending on the action taken at X. We replace the original game with the modified game of Fig. 2, in which decisions are being made at a second decision node even if the action taken in the first decision node was to exit, consistent with the procedure employed by Deck and Sarangi (2009). Thus, in the modified game, the decision maker always passes through the information set exactly twice. However, if the decision maker exits at the first node, the payoff is the same as in the original game regardless of the action taken at the second node. In this section, we analyze the modified game following the analysis of PR and AHP.

The behavioral strategy is defined by the probability of choosing “continue” at the information set, denoted by p . The expected payoff at the planning stage is given by

$$(1 - p)a + p(1 - p)b + p^2c.$$

The planning-optimal strategy is therefore

$$p^* = \frac{b - a}{2(b - c)}. \quad (1)$$

When reaching the information set in the action stage, the decision maker forms beliefs about the decision node he is in. Denote by α_X , α_Y , and $\alpha_{Y'}$ the probabilities that the decision maker assigns to being at X , Y , and Y' , respectively. In line with PR’s analysis, the expected payoff is given by

$$\alpha_X[(1 - p)a + p(1 - p)b + p^2c] + \alpha_Y[(1 - p)b + pc] + \alpha_{Y'}a. \quad (2)$$

As noted above, since the driver passes through X and through either Y or Y' exactly once, any consistent belief must be such that $\alpha_X = \alpha_Y + \alpha_{Y'} = 0.5$. Maximizing (2) over p therefore yields the action-optimal strategy

$$\bar{p} = \frac{b - a}{2(b - c)} - \alpha_Y, \quad (3)$$

which is strictly smaller than p^* for any $\alpha_Y > 0$.

AHP claim that the above analysis of action-optimal strategy is “flawed” (p. 102) in its formulation. They argue for two normative requirements. First, “when at one intersection, [the decision maker] can determine the action *only there* and *not* at the other intersection” (emphasis in original). Second, “whatever reasoning obtains at one [intersection] must obtain also at the other.” Denote by q the behavioral strategy that the decision maker who reaches the information set expects to play at the other intersection (equivalent to the strategy of the “twin-self” in PR’s multiself approach). The expected payoff to be maximized is now

$$\alpha_X[(1 - p)a + p(1 - q)b + pqc] + \alpha_Y[(1 - p)b + pc] + \alpha_{Y'}a. \quad (4)$$

Substitute the consistent beliefs $0.5q$ and $0.5(1 - q)$ for α_Y and $\alpha_{Y'}$, respectively, and let q be the planning-optimal strategy derived in (1). It is easy to confirm that (4) reduces to

$$\frac{a}{2} + \frac{b(b - a)}{4(b - c)}, \quad (5)$$

which does not depend on p . Hence p^* maximizes (5) and is both planning-optimal and action-optimal. Generally, any planning-optimal strategy must be action-optimal (Proposition 3 of PR), although the opposite is not generally true (Section 5 of AHP).

3 Experimental design

Deck and Sarangi (2009) have shown how absentmindedness can be induced in the lab by means of cognitive overload resulting from information abundance.⁸ In their experiment, participants were presented with a game tree in each period and had five seconds to choose an action. In addition to the strict time constraint, cognitive load was enhanced by including a simple matching task in each period. In the action stage, each game tree was presented exactly twice, regardless of the choice made in its first appearance. The actions chosen in the two occurrences were implemented in order of appearance to determine payoffs. The time constraint and multiple tasks are assumed to impair memory performance (see, e.g., the divided attention model of Kahneman 1973). Participants are thus not likely to recognize whether any occurrence of a specific game tree is the first or the second. As an indirect manipulation check, Deck and Sarangi (2009) verified, in a second experimental stage, that participants were unable to correctly identify the game trees included in the action stage.

Following the procedure introduced by Deck and Sarangi (2009), we implemented an environment in which many driving “maps” are presented to the participants. We adapted the original procedure in several ways. First, an abundance of evidence has shown that time pressure leads to intuitive decisions, whereas we are interested in testing whether time inconsistencies arise in deliberate decisions, which are more likely to conform to fully rational analysis (Kahneman 2011). Therefore we did not impose any time constraint. Second, to increase the difficulty of recalling specific maps, each game tree (defined by the payoffs a , b , and c), appeared in four distinct maps differing in background color (yellow, green, blue, or purple). Furthermore, we provided a direct manipulation check by explicitly asking participants in half of our treatments to guess at which decision node they are (first or second).

We constructed five different treatments: a planning-stage treatment and four action-stage treatments. The variations between the four action-stage treatments will be fully described in Sect. 3.2. To provide a crisp comparison of a planning and an action stage, we employed a within-subjects design in which each participant was exposed to both the planning-stage treatment and one of the four action-stage treatments. Participants chose behavioral strategies for the same maps in both treatments. The planning-stage and action-stage treatments differed in that one choice of behavioral strategy per map was elicited in the planning stage, whereas two choices per map were elicited in the action stage. Thus, when called upon to make an action-stage (but not a planning-stage) decision, participants could form beliefs about which node their current choice applies to. The following subsections elaborate on the implementation of the two stages. A translation of the full German instructions is available in the Supplementary materials.

⁸ When people must process large amounts of information within a short time span, the limited capacity of their short-term memory causes cognitive overload (see, e.g., Kareev and Warglien 2003). Short-term memory capacity refers to the number of items that an individual can retain at one time and is classically estimated to be 7 ± 2 (Miller 1956; Shiffrin 1976; Kareev 2000); but see Cowan (2001), for a lower estimate.

3.1 Planning stage

Participants made planning-stage decisions for a total of 56 maps constructed by crossing 14 game trees of the type depicted in Fig. 1 with the four background colors. To elicit (and allow for) behavioral strategies, we employed a mechanism similar to that used by [Huck and Müller \(2002\)](#).⁹

Participants were asked to imagine an urn with 100 balls, each marked with either “exit” or “continue”. They could determine the composition of the urn by choosing the numbers of “exit” and “continue” balls. Their choices were restricted to integer values between 0 and 100, and had to sum to 100. Once the composition of the urn was decided, the computer randomly drew (with replacement) one ball from the urn. A draw of an “exit” ball implied taking the first exit for a payoff of a . A draw of a “continue” ball implied continuing to the second intersection and drawing a second ball. If this was an “exit” ball, then the participant took the second exit for a payoff of b ; otherwise, he obtained a payoff of c . We made clear that the *same* behavioral strategy (as chosen by the participant) would be implemented in both intersections.

Participants could familiarize themselves with the task—including both the game form and the randomization mechanism—in a preliminary 15-minutes practice stage using the following procedure. Participants filled in three payoffs in a blank tree and chose a behavioral strategy. After clicking a confirm button, the resulting probability distributions over the three payoffs and the expected payoff (based on the chosen strategy) were presented on screen. Next, particular payoff realizations could be obtained by pressing a “travel”-button repeatedly.

3.2 Action stage

The action stage commenced after the completion of the planning stage. The procedure followed that of the planning stage, with the main exception that each map appeared exactly twice within the stage. As noted above, there were four different treatment variations in the action stage. We start by describing the basic treatment, followed by a discussion of its variations.

The behavioral strategy chosen in the first encounter applied to the first decision node (X in Fig. 2). The behavioral strategy chosen in the second encounter of the same map applied to the second decision node (Y and Y' in Fig. 2).¹⁰ The same map was never shown in two consecutive periods.

Sixteen *filler maps* (4 trees \times 4 colors) were added to the 56 experimental maps for a total of 72 maps or 144 game decisions (experimental periods). The filler maps served the purpose of filling up the first ten periods when participants can be certain to see a map for the first time. Additionally, initial maps are likely to be recalled later due to the well-studied primacy effect ([Murdock 1962](#); [Deese and Kaufman 1957](#)).

⁹ Recall that eliciting pure actions in this stage, as done by [Deck and Sarangi \(2009\)](#), effectively implements mixed rather than behavioral strategies.

¹⁰ That is, in line with the analysis of Sect. 2 and the procedure used in [Deck and Sarangi \(2009\)](#), but unlike the original game of Fig. 1, assigning a probability of 0 to continue in the first encounter of a map did not preclude encountering it again.

We term this procedure “induced absentmindedness without belief elicitation” (henceforth IND-WITHOUT treatment). We complemented it with three additional treatments created by manipulating two variables in a 2×2 between-subjects design as explained in the following.

3.2.1 Induced versus imposed absentmindedness

In our baseline IND-WITHOUT treatment, the first decision node always preceded the second decision node of the same map. The participants could utilize this information to shape their strategies. Specifically, they could choose “continue” with high probability in earlier periods, which were likely to correspond to the first node X , reducing this probability in later periods of the experiment.

Hence, for a robustness check, we added a further treatment that we term “imposed absentmindedness” (henceforth IMP treatment). In this treatment, the probabilities chosen by the participants in the two occurrences of any map were randomly matched to decision nodes. Either the natural order (X before Y/Y') or the reverse order (Y/Y' before X) was implemented with equal probabilities.

Under this procedure, a participant who faced an action-stage map knew that with probability 0.5 his current decision would apply to decision node X , with some probability $\alpha_Y \leq 0.5$ it would apply to decision node Y , and with the remaining probability it would not affect his payoff, as in decision node Y' . Thus, this procedure ensures that the consistent belief assigned to X is 0.5, as assumed in the theoretical analysis. Other possible interpretations exist for this procedure.¹¹ While the interpretation we provide above appears more likely to reflect the participants’ perceptions of the decision problem, it is not necessary to commit to it when testing for time inconsistencies, since the theoretical approach that does not entail time consistencies (PR’s multiself or AHP’s approach) is robust to the different interpretations.

In some respects, the IMP treatment could have been conducted by simply asking participants for two behavioral strategies per map. While this avoids the extended procedure involving memory overload, it would open the door to a different kind of strategy: participants could coordinate their decisions by choosing “continue” at one node and “exit” at the other. This would guarantee them an expected payoff of $\frac{a+b}{2}$, which is higher than the payoff obtainable by any behavioral strategy in all of our experimental trees (albeit not generally, see Section 6(e) of AHP for further discussion). This is equivalent to the asymmetric equilibria of the game that results from letting two players play the role of the decision maker, with the decision of each player randomly assigned to one of the two decision nodes (Gilboa 1997). Huck and Müller (2002) found experimentally that players learn to coordinate on such an equilibrium after repeated play. Memory overload hampers the possibility of carrying

¹¹ One is that it gives rise to a game of absentmindedness with six decision nodes in the information set: three corresponding to the ones in Fig. 2 and three for the reverse-order game. Another is that the game form is the same as in Fig. 2, with the final payoffs being the expected payoffs over the two orders. While the planning-optimal strategy remains the same in all interpretations, the action-optimal strategy can vary substantially. In particular, the paradox does not necessarily hold for some (inconsistent) beliefs.

Table 1 Measurement of participants' confidence in their guesses

Your choice	Option	If your guess is <i>correct</i>	If your guess is <i>wrong</i>
		you WIN	you LOSE
o	A	1	1
o	B	3	5
o	C	5	15

out such a strategy that requires the identification of different occurrences of the same map.¹²

3.2.2 Belief elicitation

In half of our sessions, we elicited (i) beliefs about the current decision node and (ii) confidence in one's own beliefs. The reason for introducing this design feature is twofold. First, it allows us to check for the actual occurrence of absentmindedness in the IND treatment. Second, since in the IMP treatment stating beliefs is tantamount to guessing the outcome of a fair coin toss, comparing beliefs and bets in this treatment with those in the IND treatment allows us to assess the *subjective feeling* of absentmindedness in the latter.

The elicitation procedure was as follows. In each period, participants were asked to guess whether they were at the first or second decision node and to place a bet on their guess being correct. We measured participants' confidence in their guesses by letting them choose one of the three options depicted in Table 1. Each option is associated with a gain and a loss depending on the guess being, respectively, correct or incorrect. The possibility of gains should incentivize participants to remember maps, even though the concomitant possibility of losses should urge those who suffer from imperfect recall to select option A.¹³

Both IND and IMP treatments were conducted without and with eliciting beliefs to control for the effect of belief elicitation. The resulting four action-stage treatments are IND-WITHOUT, IMP-WITHOUT, IND-WITH, and IMP-WITH.

3.3 Experimental game trees

The game trees used in the experiment are shown in Table 2. Trees 1–14 were presented to the participants in both the planning stage and the action stage. Trees 15–18 were

¹² Our data confirm that participants in the IMP treatment did not follow this strategy, which would drive the mean probability assigned to "continue" in the action stage towards 0.5.

¹³ The numbers in Table 1 are chosen so that the expected payoff from option A exceeds the expected payoff from the other two options whenever the probability assigned to being correct is lower than 2/3. Only if the probability of being correct is greater than 5/6, a risk neutral decision maker should opt for C. In principle, the belief elicitation procedure gives participants an opportunity to diversify since they can play according to their true beliefs and state the opposite beliefs. Our results, however, show that participants actually play in line with their stated beliefs.

Table 2 Experimental game trees

Tree number	<i>a</i>	<i>b</i>	<i>c</i>	<i>p</i> *	$\alpha = 1/(1 + p^*)$	Tree type
1	20	50	30	0.75	0.57	Paradox
2	10	80	30	0.70	0.59	Paradox
3	0	40	10	0.67	0.60	Paradox
4	10	50	20	0.67	0.60	Paradox
5	30	90	40	0.60	0.63	Paradox
6	30	70	30	0.50	0.67	Paradox
7	20	80	10	0.43	0.70	Paradox
8	30	60	20	0.38	0.73	Paradox
9	30	70	10	0.33	0.75	Paradox
10	30	50	10	0.25	0.80	Paradox
11	50	10	30			Optimal exit
12	60	40	20			Optimal exit
13	20	10	50			Optimal continue
14	10	30	60			Optimal continue
15	0	50	10			Filler
16	10	80	40			Filler
17	30	90	50			Filler
18	10	60	20			Filler

used to construct the filler maps that only appeared in the action stage and are excluded from the analysis.

In addition to ten *paradox trees*, we included two *optimal exit trees* (in which $a > b, c$) and two *optimal continue trees* (in which $c > a, b$). The planning-optimal, as well as the action-optimal, behavioral strategy in the optimal exit and optimal continue trees is to assign a probability of, respectively, 0 and 1 to continue. Hence, no time inconsistencies are expected. Behavior in these trees will provide a check of the participants' understanding of the task.

3.4 Procedures

The computerized experiment was conducted in the experimental laboratory of the Max Planck Institute of Economics (Jena, Germany) in April and December 2009. It was programmed in z-Tree (Fischbacher 2007). The participants were undergraduate students from the Friedrich Schiller University of Jena. They were recruited using the ORSEE software (Greiner 2004). Upon entering the laboratory, the participants were randomly assigned to visually isolated computer terminals. The initial instructions stated that the experiment consisted of two phases, and only included the rules of the first phase (corresponding to the planning stage). Written instructions on the second phase (corresponding to the action stage) were distributed after the first phase was completed. Before starting the experiment, participants had to answer a control questionnaire testing their comprehension of the rules.

We ran eight experimental sessions, two for each of the four between-subjects treatments. Each session included 18 participants so that we have 36 participants per between-subjects treatment, and 144 participants overall. In all sessions, participants

did not know the number of maps (and thus periods) beforehand. Moreover, to avoid wealth effects, they did not receive any feedback about the random draws determining their period payoff (i.e., the realization of their strategy) and their earnings from the guesses until the end of the session.

Each session lasted about two hours. Money in the experiment was denoted in ECU (Experimental Currency Unit), where 10 ECU = €0.07. Final payoff was the sum of all payoffs accumulated over the two stages.¹⁴ The average earnings per participant were €35.40 (including a €2.50 show-up fee).

4 Experimental results

Before testing our main hypothesis that participants exit with higher probability in the action stage than in the planning stage, we test the crucial auxiliary hypotheses. We start by looking at the relation between choices and optimal strategies in the optimal exit and optimal continue trees to check that participants understood the task and behaved in line with the incentives. Next, we analyze behavior in the action stage to verify that participants were indeed absentminded both objectively (as reflected by the correlation between decisions and actual nodes) and subjectively (as reflected by the elicited beliefs and bets). We conclude this section with the empirical test of the main hypothesis.

4.1 Implementation of the game

For the optimal exit and optimal continue trees, we expect participants to behave optimally in close to 100 % of the decisions. Table 3 shows that the majority of choices for these trees conform to the optimal strategy in both stages. In the planning stage, the proportion of optimal choices for trees 11–14 is above 90 % and the mean strategy is to take the optimal action with over 0.95 probability.¹⁵ Although the proportion of optimal choices is lower in the action stage, the mean strategies are close to optimal in all but one case. The apparently odd behavior observed for tree 11 can be explained by participants choosing as if they knew which node they are at. To see this, consider the optimal strategies that would result from informing the participants about the actual node. Formally, change the game in Fig. 2 to have two information sets, one including node X and the other including nodes Y and Y' .¹⁶ In trees 12–14, the optimal strategies in the two new information sets coincide, as in the original game. Conversely, in tree 11, the optimal strategies in the two information sets differ. We conjecture that participants are sometimes swayed by the possibility of being at decision node Y . In tree 11, this makes it optimal to continue (for a payoff of 30 rather than 10) even though the optimal behavioral strategy in this tree is to exit.

¹⁴ By paying a small monetary amount over a large number of periods we try to induce risk neutrality.

¹⁵ As the planning-stage treatment is identical for all participants and the participants are unaware of the following action stage, we pool the planning-stage data. We average choices for each participant over the four instances of each tree (defined by background colors) to increase reliability.

¹⁶ This turns the game into a game with imperfect recall, but not absentmindedness.

Table 3 Continue choices for the optimal exit (11,12) and optimal continue (13,14) trees

Tree	Proportion of optimal choices		Mean strategy ^a	
	Planning (%)	Action (%)	Planning	Action
11	92.01	56.16	2.07 (0.51)	22.56 (2.13)
12	93.58	77.60	1.99 (0.55)	6.86 (1.04)
13	92.36	84.03	96.78 (0.90)	94.77 (0.91)
14	92.88	87.59	97.50 (0.67)	96.14 (0.75)

^a Standard errors (based on 144 observations) in parentheses

Turning to the paradox trees, we do not expect participants to be able to compute the exact optimal strategy p^* reported in Table 2 and derived from Eq. (1). However, if participants are reacting correctly to the payoffs they can obtain from each tree, their choices should be correlated with the optimal strategy across trees. Averaging over the 144 participants for each of the 10 paradox trees, we find that choices in these trees are strongly and significantly correlated with p^* ($r = 0.906$, $p < 0.001$). Thus, we conclude that the participants managed to cope with the complexity of the task and understood the experimental incentives.

4.2 Implementation of absentmindedness

If participants in the IND treatments were able to recall that they had seen a map in a previous period, we would expect them to choose a higher probability to continue in the first node. The random-effects GLS models reported in the top panel of Table 4, with the probability assigned to continue as the dependent variable and the actual node as our main independent variable, reveal a weak effect in this direction. Model 1, controlling for the theoretical p^* , establishes that, on average, the probability to continue is roughly 0.07 higher in the first node than in the second one. In the IND-WITHOUT treatment, the effect is lower, just above 0.03 (see Model 5).

Nonetheless, the effect of the actual node could be due to the correspondence between actual node and period that rises from the experimental procedure. Recall that participants in the IND treatments knew that, for each map, the first node always appeared before the second one. Taking this into account, they could improve their expected payoff by choosing to continue with higher probability at the beginning of the session. Models 2 and 6 in Table 4 show that the coefficient of *Period* is significantly negative, implying that participants indeed tend to decrease the probability to continue as the experiment progresses in both IND treatments. Additionally, once the period is included, the effect of the actual node drops from 0.07 to less than 0.03 in the

Table 4 Generalized least-squares random-effects regressions on action-stage strategies

	IND-WITH				IND-WITHOUT	
	(1)	(2)	(3)	(4)	(5)	(6)
Intercept	27.825*** (3.370)	27.814*** (3.340)	83.167*** (2.904)	82.702*** (2.918)	21.265*** (2.760)	21.262*** (2.777)
<i>Actual node</i>	-6.901*** (0.999)	-2.482** (1.258)	-2.251*** (0.823)	-0.365 (1.030)	-3.231*** (0.892)	-2.014* (1.128)
<i>p</i> *	0.353*** (0.030)	0.378*** (0.030)	0.382*** (0.025)	0.392*** (0.025)	0.472*** (0.027)	0.479*** (0.027)
<i>Belief</i>			-37.411*** (0.971)	-37.099*** (0.975)		
<i>Period</i>		-0.110*** (0.019)		-0.048*** (0.016)		-0.303* (0.017)
Observations	2,880	2,880	2,880	2,880	2,880	2,880
<i>N</i> participants	36	36	36	36	36	36

	IMP-WITH				IMP-WITHOUT	
	(7)	(8)	(9)	(10)	(11)	(12)
Intercept	7.387*** (2.563)	8.906*** (2.586)	22.270*** (2.779)	23.225*** (2.783)	5.744** (2.492)	11.445*** (2.588)
<i>Actual node</i>	-0.690 (0.725)	-0.676 (0.725)	-0.589 (0.703)	-0.580 (0.703)	-0.249 (0.726)	-0.232 (0.717)
<i>p</i> *	0.770*** (0.022)	0.775*** (0.022)	0.776*** (0.021)	0.780*** (0.021)	0.804*** (0.022)	0.825*** (0.026)
<i>Belief</i>			-11.202*** (0.822)	-11.121*** (0.823)	(0.745)	(0.746)
<i>Period</i>		-0.026** (0.011)		-0.018* (0.011)		-0.093* (0.011)
Observations	2,880	2,880	2,880	2,880	2,880	2,880
<i>N</i> participants	36	36	36	36	36	36

Note. Numbers in parentheses are estimated standard errors.

*,**,*** Significance at the 10, 5, and 1 %, respectively

IND-WITH treatment and from around 0.03 to 0.02 in the IND-WITHOUT treatment. In the latter case, the coefficient of *Actual node* is only significant at the 10 % level.

In the IND-WITH treatment, an additional control is provided by including the beliefs (Models 3 and 4 in Table 4) because participants are likely to base their beliefs on the period in the same way as they base their game decisions on it. Including the beliefs in the model has a similar effect to including the period, and once both independent variables are taken into account, the effect of the actual node, reflecting true recall, completely disappears. Although this is not a clean control (as the beliefs could also indicate lack of absentmindedness), the limited effect size of *Actual node* in Models 2 and 6 as well as the statistical insignificance of its coefficient in the full Model 4 suggest that true recall is strongly limited. Furthermore, true recall is ruled

Table 5 Beliefs and bets in IND-WITH

Bet	Overall (%)	By belief	
		X	Y
		62.8 %	37.2 %
A	85.2	87.1 %	84.4 %
B	9.7	10.2 %	9.5 %
C	5.1	2.7 %	6.1 %

out by design in the IMP treatments, as reflected in Models 7–12. These treatments can thus serve to complement the analysis conducted on the IND treatments in the next section.

We proceed to look at the bets made on the beliefs in the IND-WITH treatment, which we use as a proxy for the *subjective* absentmindedness experienced by the participants. The safest bet *A* is the optimal bet not only for risk-neutral and absentminded participants, but also for less-than-absentminded participants who are sufficiently risk-averse. To control for risk attitudes, we compare bets made in the IND-WITH treatment to bets made in the IMP-WITH treatment, where participants lack any knowledge of the actual node they are at, so that only risk attitudes can affect their betting choices.

In the IMP-WITH treatment, the safest bet *A* (indicating lack of confidence in the beliefs) was chosen 80.7 % of the time. The riskier bets *B* and *C* were chosen 18.4 and 1.0 % of the time, respectively. In comparison, participants in the IND-WITH treatment chose bets *A* and *C* more often (85.2 and 5.1 %, respectively), while they choose the intermediate bet *B* less frequently (9.7 %). As can be seen in Table 5, betting behavior was similar regardless on which decision node participants guessed to be at. Thus, participants were not systematically more confident about their bets in the IND-WITH treatment than they were in the IMP-WITH treatment, indicating that they indeed felt absentminded.

To conclude, on the whole, participants in the IND treatments were and felt mostly, though not perfectly, absentminded. The degree to which they were able to recall the previous occurrence of a current map is estimated at approximately 2–3 %, reflecting the coefficient of *Actual node* in Models 2 and 6 in Table 4 as well as the increase in the risky bets *C* compared to in the IMP-WITH treatment. To account for this small degree of true recall, we test the robustness of the results reported in the next section using two measures. First, by analysing participants' decisions in the periods in which participants indicated low confidence in their memory by choosing bet *A*. Second, by analysing behavior in the IMP treatments, which are free from issues of induced absentmindedness.

4.3 Planning versus action

We turn now to the main hypothesis. Did our participants exit with higher probability in the action stage compared to the planning stage? The mean strategies by paradox tree and stage are summarized in Tables 6 and 7. A graphical representation of the same data is provided in Fig. 3, where the trees are ordered on the horizontal axis by $\sim p^*$.

Table 6 Continue choices for the paradox trees in the IND treatments

Tree	p^*	IND-WITH				IND-WITHOUT				
		Plan ^a (1)	Action ^a (2)	Plan versus action p -value ^b (3)	Action $\beta = X^a$ (4)	Action $\beta = Y^a$ (5)	$\beta = X$ versus $\beta = Y$ p -value ^c (6)	Plan ^a (7)	Action ^a (8)	Plan versus action p -value ^b (9)
1	75.000 (2.849)	62.951 (2.849)	42.278 (3.624)	0.000	73.504 (3.328)	36.391 (3.908)	0.000	63.743 (1.601)	49.701 (2.589)	0.000
2	70.000 (3.121)	67.799 (3.121)	41.806 (3.714)	0.000	70.333 (4.016)	29.515 (3.558)	0.000	68.167 (2.149)	50.660 (2.679)	0.000
3	66.667 (3.363)	77.965 (3.363)	44.410 (4.242)	0.000	84.336 (3.701)	31.732 (4.505)	0.000	76.646 (2.004)	53.448 (3.531)	0.000
4	66.667 (3.218)	65.993 (3.218)	40.257 (4.132)	0.000	64.878 (4.789)	32.393 (4.071)	0.000	61.875 (1.542)	48.274 (2.955)	0.000
5	60.000 (3.142)	63.278 (3.142)	39.899 (3.843)	0.000	70.310 (2.949)	25.583 (3.779)	0.000	60.507 (2.049)	43.549 (3.027)	0.000
6	50.000 (2.332)	52.903 (2.332)	35.139 (2.936)	0.000	69.301 (3.732)	19.345 (2.722)	0.000	53.750 (1.780)	35.833 (2.582)	0.000
7	42.857 (3.169)	44.146 (3.169)	29.948 (3.000)	0.001	54.356 (4.419)	17.719 (2.976)	0.000	48.889 (1.819)	33.198 (2.272)	0.000
8	37.500 (3.180)	46.222 (3.180)	30.101 (3.002)	0.000	52.313 (3.958)	21.610 (2.461)	0.000	45.174 (2.223)	35.122 (2.387)	0.000
9	33.333 (3.083)	41.604 (3.083)	29.385 (3.048)	0.000	51.665 (4.224)	19.611 (2.678)	0.000	46.215 (2.351)	32.528 (2.779)	0.000
10	25.000 (3.037)	40.382 (3.037)	27.556 (2.490)	0.000	45.529 (4.302)	20.632 (2.731)	0.000	45.139 (2.272)	30.750 (2.661)	0.000

^a Standard errors (based on 36 independent observations) in parentheses

^b Two-sided Wilcoxon signed rank tests with continuity correction relying on 36 independent observations (averages over the 4 colors for each participant)

^c Two-sided Wilcoxon signed rank tests with continuity correction. Number of independent observations for each map: $N_1 = 31$; $N_2 = 26$; $N_3 = N_4 = 32$; $N_5 = 31$; $N_6 = 29$; $N_7 = 27$; $N_8 = N_9 = 30$; $N_{10} = 33$

The mean strategies are lower in the action stage than in the planning stage for all 10 paradox trees and in all four treatments (compare column (1) with (2) and column (7) with (8) in Tables 6 and 7).¹⁷ This difference is statistically significant in all but 3 cases according to Wilcoxon signed rank tests with continuity correction relying on 36 independent observations (see columns (3) and (9) in the tables). Overall, the probability assigned to continue in the action stage is, on average, 69.0% (85.1%) of that assigned in the planning stage, in the IND (IMP) experiment.

Columns (4) and (5) in Tables 6 and 7 list the mean strategies by beliefs (β) elicited in the WITH treatments. The comparison of the two columns reveals that participants' strategies are strongly aligned with their beliefs (namely participants continue with lower probability when they believe to be at Y rather than at X).¹⁸ In the IND-WITH treatment, this can (partly) reflect true recall. In the IMP-WITH treatment, however, beliefs are mere guesses, so that payoff maximization cannot explain the observed contingency. This result supports our interpretation of the observed behavior in the optimal-exit game tree 11, i.e., that participants sometimes behave as if they were not absentminded.

To complement the non-parametric tests, we computed the planning-action gap to use as a dependent variable in a regression analysis. The planning-action gap is defined as the difference between the probability that a participant assigned to continue in an action-stage decision and the average probability that he assigned to continue over the four instances of the same game tree that appeared in the planning stage. A negative gap indicates that the participant chose to exit with a higher probability in the action stage.

Table 8 reports the results of a series of regressions on the planning-action gap. The models in columns (1)–(3) include the treatment variables IMP and WITH and the planning-optimal strategy p^* , gradually adding its second and third powers. Since the IND and IMP treatments employ inherently different procedures, columns (4) and (5) add the interactions of the procedure with the other independent variables. Columns (6) and (7) replace the belief-elicitation variable WITH by the three levels of confidence A , B , and C and their interactions with the procedure, taking the WITHOUT treatment as the baseline.

The results of the regressions can be best understood by looking at Fig. 4. The figure plots the predicted planning-action gap according to the full (and most comprehensive) model of column (7) against the different p^* 's. Several conclusions can be drawn from the figure and the accompanying regressions. First, the predicted gap is negative across the four treatments and range of game-trees, reflecting the paradox as predicted by PR's analysis. Second, belief elicitation does not appear to affect behavior in any substantial way, as reflected in the insignificant coefficient of WITH in models (1)–(5). Third, the higher the confidence in one's own beliefs, the larger the planning-action gap. Since in the IND-WITH treatment this may reflect true recall, it is important to note that, even for low-confidence decisions, the planning-action gap is far from being negligible. The small degree of true recall apparent in the data does not, therefore, provide a satisfactory

¹⁷ The connecting lines in Fig. 3 were added to visually emphasize this fact.

¹⁸ Note that diversification in order to reduce risk implies the opposite.

Table 7 Continue choices for the paradox trees in the IMP treatments

Tree	p^*	IMP-WITH				IMP-WITHOUT				
		Plan ^a (1)	Action ^a (2)	Plan versus action p -value ^b (3)	Action $\beta = X^a$ (4)	Action $\beta = Y^a$ (5)	$\beta = X$ versus $\beta = Y$ p -value ^c (6)	Plan ^a (7)	Action ^a (8)	Plan versus action p -value ^b (9)
1	75.000	61.458 (2.138)	56.361 (2.504)	0.036	60.588 (2.732)	46.019 (3.643)	0.001	63.882 (2.175)	58.427 (2.226)	0.085
2	70.000	63.181 (2.859)	61.844 (2.541)	0.278	66.745 (2.979)	48.903 (4.249)	0.003	67.840 (2.267)	60.601 (2.605)	0.023
3	66.667	74.076 (2.941)	69.354 (2.464)	0.006	74.762 (2.449)	59.426 (3.926)	0.001	80.007 (2.078)	76.177 (2.796)	0.221
4	66.667	60.722 (1.986)	59.132 (2.832)	0.391	62.781 (2.872)	49.559 (3.931)	0.001	63.840 (2.263)	58.028 (2.393)	0.066
5	60.000	56.667 (2.066)	51.302 (3.076)	0.047	56.205 (3.446)	39.809 (3.863)	0.001	61.028 (2.540)	50.962 (2.630)	0.002
6	50.000	50.653 (1.871)	41.545 (2.682)	0.001	46.325 (2.978)	34.473 (3.626)	0.000	53.250 (1.722)	43.059 (2.489)	0.001
7	42.857	43.174 (2.277)	34.556 (2.903)	0.003	37.828 (3.057)	31.256 (4.163)	0.016	48.035 (2.936)	35.385 (3.334)	0.000
8	37.500	42.389 (2.323)	34.194 (2.770)	0.000	36.844 (2.880)	30.222 (3.968)	0.005	43.778 (2.834)	34.604 (2.720)	0.000
9	33.333	39.215 (2.692)	30.299 (2.932)	0.000	31.071 (3.021)	29.275 (3.615)	0.024	42.917 (2.835)	30.375 (2.956)	0.000
10	25.000	36.993 (2.644)	30.490 (2.611)	0.020	32.046 (2.995)	30.546 (3.548)	0.067	40.875 (2.402)	30.028 (2.621)	0.000

^a Standard errors (based on 36 independent observations) in parentheses^b Two-sided Wilcoxon signed rank tests with continuity correction relying on 36 independent observations (averages over the four colors for each participant)^c Two-sided Wilcoxon signed rank tests with continuity correction. Number of independent observations for each map: $N_1 = 25$; $N_2 = N_3 = 28$; $N_m = 27$ for $m = 4, \dots, 9$; $N_{10} = 28$

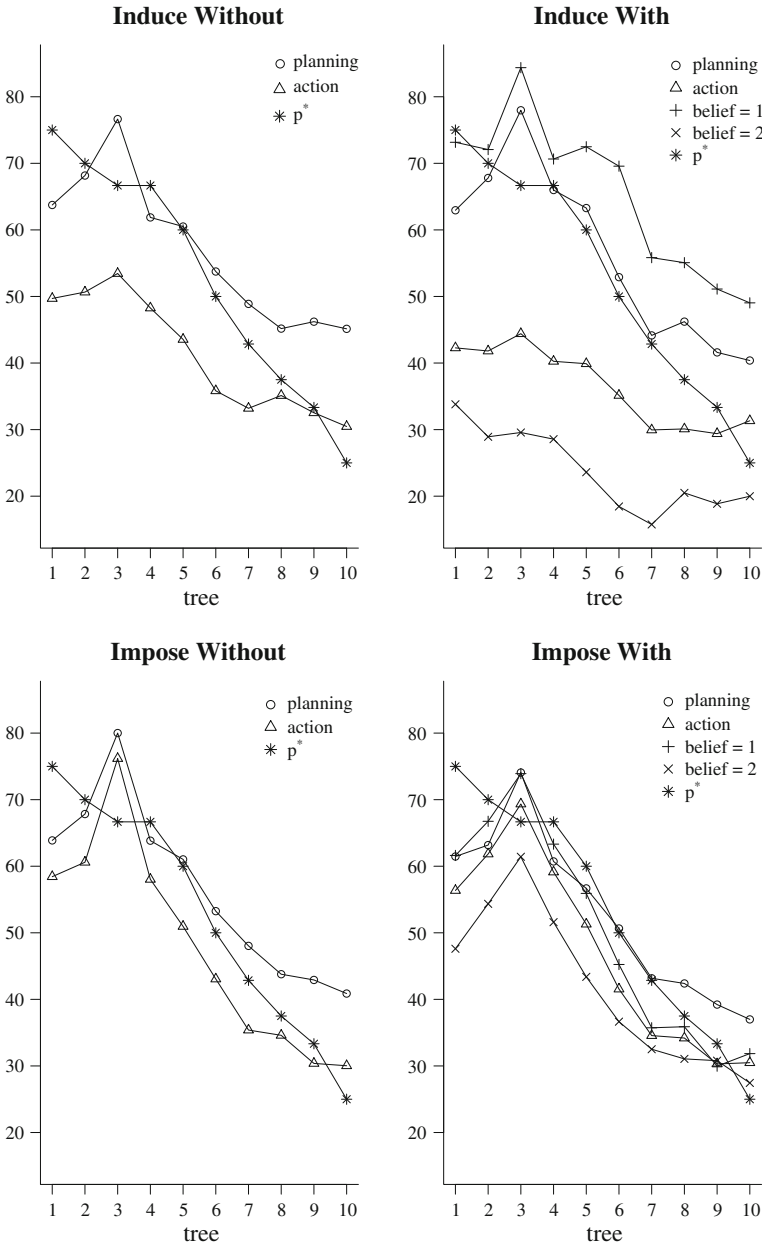


Fig. 3 Mean continue choices in the four treatments

Table 8 Generalized least-squares random-effects regressions on the planning-action gap

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Intercept	-15.53*** (2.048)	-7.792** (3.133)	-29.05*** (7.575)	-66.46*** (10.44)	-64.61*** (10.50)	-65.96*** (10.41)	-63.92*** (10.44)
IMP	10.61*** (2.229)	10.61*** (2.229)	10.61*** (2.244)	85.43*** (14.68)	81.74*** (14.85)	86.05*** (14.63)	81.90*** (14.77)
WITH	-0.853 (2.229)	-0.853 (2.229)	-0.853 (2.244)	-0.853 (2.237)	-4.542 (3.144)		
IMP × WITH					7.378* (4.447)		
$\frac{p^*}{100}$	-3.823*** (1.294)	-38.04*** (10.56)	108.9** (48.81)	407.4*** (68.39)	407.4*** (68.39)	396.1*** (68.21)	393.2*** (68.28)
$\frac{p^*}{100}^2$		33.75*** (10.34)	-281.2*** (102.7)	-948.2*** (143.9)	-948.2*** (143.9)	-908.2*** (143.5)	-899.8*** (143.7)
$\frac{p^*}{100}^3$			211.7*** (68.67)	655.0*** (96.22)	655.0*** (96.22)	619.6*** (96.03)	612.7*** (96.18)
IMP × $\frac{p^*}{100}$				-597.0*** (96.72)	-597.0*** (96.72)	-593.1*** (96.43)	-588.2*** (96.56)
IMP × $\frac{p^*}{100}^2$				1,334*** (203.5)	1,334*** (203.5)	1,309*** (202.9)	1,296*** (203.2)
IMP × $\frac{p^*}{100}^3$				-886.6*** (136.1)	-886.6*** (136.1)	-861.0*** (135.7)	-851.5*** (135.9)
A						0.367 (2.222)	-2.960 (2.889)
B						-4.890** (2.460)	-10.36*** (3.221)
C						-15.58*** (2.887)	-19.98*** (3.521)
IMP × A							6.303 (4.098)
IMP × B							11.64** (4.618)
IMP × C							10.09 (6.324)
Observations	11,520	11,520	11,520	11,520	11,520	11,520	11,520
N participants	144	144	144	144	144	144	144

Note. Numbers in parentheses are estimated standard errors
 *, **, *** Significance at the 10, 5, and 1 %, respectively

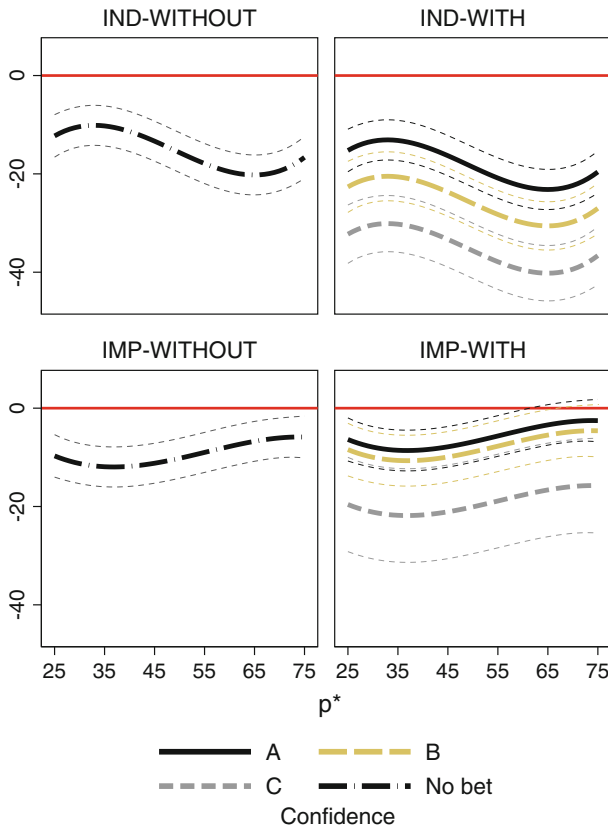


Fig. 4 Planning-action gap for the four between-subjects treatments with 95 % confidence intervals. In the WITH treatments, the gap is displayed by confidence level

explanation for the observed time inconsistencies. In the IMP-WITH treatment, on the other hand, bets are pure gambles and, as such, should be independent of the game choices. The apparent deviation from this principle suggests that, even when participants are exogenously forced to be absentminded, they sometimes feel as if they knew which decision node they are at. This is reflected in both the willingness to take the risky bet C and, when it is indeed taken, to assign a significantly higher probability to exit.

5 Conclusion

Imperfect recall surrounds us wherever we go. We often fail to recall where we left our keys or parked our car. At times we struggle to remember whether we actually got around to paying that pending electricity bill or registering for the upcoming conference. The situation might be even worse when we think of decisions taken by organizations, such as firms or countries, where frequently ‘the right hand doesn’t know what the left hand is doing’. Nonetheless, the vast majority of theoretical and

experimental research conducted to understand rational decision making has so far been confined to situations of perfect recall.

Some important issues arising from imperfect recall are well illustrated by the paradox of the absentminded driver.¹⁹ This paper joins and complements the theoretical efforts devoted to the paradox by providing a positive analysis of the problem. Specifically, we report on an experiment designed to compare behavior in a planning stage and an action stage of a decision problem featuring absentmindedness. We find that participants tend to exit more in the action stage than in the planning stage. This planning-action gap is robust to different parameterizations of the problem, and is apparent in situations where participants' choices indicate that they feel absentminded and in an environment that guarantees absentmindedness by imposing beliefs in the information set exogenously. Overall, our data are in line with the provisional conclusion reached by [Deck and Sarangi \(2009\)](#) based on a comparison of behavior in an action stage and the theoretical benchmark, and strongly support our hypothesis as derived from PR's theoretical discussion.

Several regularities in our data are not captured by PR's analysis, and point at a behavioral principle which states that, although absentminded, people are affected by what they would have done in a similar situation that does not involve absentmindedness. In other words, people behave as if they knew the decision node they are at. These regularities are the following. First, when $a > c > b$, as in our game tree 11, the optimal strategy is to exit. Nevertheless, some participants were partly swayed by what they would have done had they known themselves to be in the second decision node, namely choose to continue. Second, in the paradox maps, participants were more likely to exit when they guessed to be in the second node. This was true even in the IMP-WITH treatment, where guesses were simple gambles. Moreover, this behavior was more pronounced when participants expressed high confidence in their guess. Again, these seemingly irrational regularities can be explained by participants sometimes behaving as if they know themselves to be in the second node.²⁰

We see several avenues for future experimental research in the realm of absentmindedness. First, additional studies are required to test the conjectured behavioral principle outlined in the previous paragraph and its relationship to PR's analysis of behavior under absentmindedness. Second, we focus on individual decision making under absentmindedness induced by information load. It would be interesting to extend our results to other manifestations of absentmindedness, and in particular absentmindedness in organizations. Finally, there are now sundry theoretical papers dealing with the fact that the expected payoff differs between the planning stage and the action stage (in addition to PR's initial effort, notable studies include [Aumann et al. 1997b](#); [Elga 2000](#); [Hitchcock 2004](#)). These provide a fertile ground for studying how people form beliefs in situations involving imperfect recall.

Acknowledgments Financial support from the Max Planck Society is gratefully acknowledged. We thank the members and students of the Center for the Study of Rationality in Jerusalem and the Max Planck Institute

¹⁹ Other directions of research focus on the analysis of players with bounded complexity (see, e.g., [Abreu and Rubinstein 1988](#); [Lehrer 1988](#); [Rubinstein 1986](#)).

²⁰ These patterns can be rationalized as risk-seeking behavior. We find this explanation to be implausible.

of Economics in Jena, and particularly Bob Aumann, Ido Erev, Werner Güth, Joe Halpern, Sergiu Hart, Yaakov Kareev, Dave Lagnado, Motty Perry, Michele Piccione, Ariel Rubinstein, Larry Samuelson, Sudipta Sarangi, Ran Shorrer, and Shmuel Zamir for helpful discussions and comments. We thank Christoph Göring for assistance with programming.

References

- Abreu D, Rubinstein A (1988) The structure of Nash equilibrium in repeated games with finite automata. *Econometrica* 56(6):1259–1281
- Aumann R, Hart S, Perry M (1997a) The absent-minded driver. *Game Econ Behav* 20:102–116
- Aumann R, Hart S, Perry M (1997b) The forgetful passenger. *Game Econ Behav* 20(1):117–120
- Bardsley N, Cubitt R, Loomes G, Moffatt P, Starmer C, Sugden R (2010) *Experimental economics: rethinking the rules*. Princeton University Press, Princeton
- Battigalli P (1997) Dynamic consistency and imperfect recall. *Game Econ Behav* 20(1):31–50
- Binmore K (1996) A note on imperfect recall. In: Albers W, Güth W, Hammerstein P, Moldovanu B, Van Damme E (eds) *Understanding strategic interaction—essays in honor of Reinhard Selten*. Springer, Berlin, pp 51–62
- Board O (2003) The not-so-absent-minded driver. *Res Econ* 57(3):189–200
- Cowan N (2001) The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behav Brain Sci* 24(1):87–114
- Deck C, Sarangi S (2009) Inducing imperfect recall in the lab. *J Econ Behav Organ* 69:64–74
- Deese J, Kaufman RA (1957) Serial effects in recall of unorganized and sequentially organized verbal material. *J Exp Psychol* 54(3):180–187
- Elga A (2000) Self-locating belief and the sleeping beauty problem. *Analysis* 60(266):143–147
- Epstein LG, Schmeidler D (2003) Recursive multiple-priors. *J Econ Theory* 113:1–31
- Fischbacher U (2007) z-tree: Zurich toolbox for ready-made economic experiments. *Exp Econ* 10(2): 171–178
- Gilboa I (1997) A comment on the absent-minded driver paradox. *Game Econ Behav* 20(1):25–30
- Gilboa I, Schmeidler D (1989) Maxmin expected utility with nonunique prior. *J Math Econ* 18:141–153
- Greiner B (2004) An online recruitment system for economic experiments. In: Kremer K, Macho V (eds) *Forschung und wissenschaftliches Rechnen 2003*. GWDG Bericht 63, Gesellschaft für Wissenschaftliche Datenverarbeitung, Göttingen, pp 79–93
- Grove AJ, Halpern JY (1997) On the expected value of games with absentmindedness. *Game Econ Behav* 20(1):51–65
- Halpern JY (1997) On ambiguities in the interpretation of game trees. *Game Econ Behav* 20(1):66–96
- Hitchcock C (2004) Beauty and the bets. *Synthese* 139(3):405–420
- Huck S, Müller W (2002) Absent-minded drivers in the lab: testing Gilboa’s model. *Int Game Theory Rev* 4(4):435–448
- Isbell J (1957) *Finitary games*. In: *Contributions to the theory of games III*. Princeton University Press, Princeton, pp 79–96
- Kahneman D (1973) *Attention and effort*. Prentice Hall, Englewood Cliffs
- Kahneman D (2011) *Thinking: fast and slow*. Farrar Straus and Giroux, New York
- Kahneman D, Tversky A (1979) Prospect theory: analysis of decision under risk. *Econometrica* 47(2): 263–292
- Kareev Y (2000) Seven (indeed, plus or minus two) and the detection of correlations. *Psychol Rev* 107(2):397–402
- Kareev Y, Warglien M (2003) Cognitive overload and the evaluation of risky alternatives: the effects of sample size, information format and attitude to risk. Discussion Paper 340, Center for the Study of Rationality
- Kline J (2005) Imperfect recall and the relationships between solution concepts in extensive games. *Econ Theory* 25(3):703–710
- Kuhn HW (1953) *Extensive games and the problem of information*. In: *Contributions to the theory of games II*. Princeton University Press, Princeton, pp 193–216
- Lehrer E (1988) Repeated games with stationary bounded recall strategies. *J Econ Theory* 46(1):130–144
- Lipman BL (1997) More absentmindedness. *Game Econ Behav* 20(1):97–101

- Miller GA (1956) The magic number 7, plus or minus 2: some limits on our capacity for processing information. *Psychol Rev* 63:81–97
- Murdock BB Jr (1962) The serial position effect of free recall. *J Exp Psychol* 64(5):482–488
- Piccione M, Rubinstein A (1997a) The absent-minded driver's paradox: synthesis and responses. *Game Econ Behav* 20(1):121–130
- Piccione M, Rubinstein A (1997b) On the interpretation of decision problems with imperfect recall. *Game Econ Behav* 20(1):3–24
- Rubinstein A (1986) Finite automata play the repeated prisoner's dilemma. *J Econ Theory* 39(1):83–96
- Shiffrin RM (1976) Capacity limitations in information processing, attention, and memory. In: Estes WK (ed) *Handbook of learning and cognitive processes*, vol 4. Erlbaum, Hillsdale, pp 177–236
- Wilson TD, Nisbett RE (1978) The accuracy of verbal reports about the effects of stimuli on evaluations and behavior. *Soc Psychol* 41:118–131