



The not-so-absent-minded driver

Oliver Board

Department of Economics, Amherst College, Amherst, MA 01002, USA

Received 1 November 2002; revised 1 March 2003

Abstract

This paper starts with a re-examination of Piccione and Rubinstein's [Games Econ. Behav. 20 (1997) 3] *Absent-Minded Driver* problem, and suggests a novel interpretation of Aumann, Hart and Perry's [Games Econ. Behav. 20 (1997) 102] notion of action-optimality. We then consider several variants of the original problem in which the assumption that the player's information sets partition the set of his decision nodes is relaxed. This relaxation enables us to construct a counter-example to Piccione and Rubinstein's result that planning-optimal strategies are always action-optimal. We also show that an agent with more information may do worse than an agent with less.

© 2003 University of Venice. Published by Elsevier Ltd. All rights reserved.

Keywords: Absent-Minded driver; Action-optimality; Decision nodes

1. Introduction

The Absent-Minded Driver problem was introduced by Piccione and Rubinstein (1997) to demonstrate an ambiguity in the interpretation of decision problems with absent-mindedness (where certain play paths pass through a given information set more than once). In particular, they observed that the decision made by a player before starting to play could exhibit time inconsistency: he would not wish to carry out his plan when called upon to do so. In a response to this paper, Aumann et al. (1997) (henceforth AHP) presented an alternative way of viewing the decision-making process of the player when on move which eliminates this time inconsistency.

It is not the aim of this paper to adjudicate between conflicting views about the right way to think about *The Absent-Minded Driver* problem. Indeed, we do not think that there is a single 'right way'. Rather in Section 2 we present one possible interpretation of the problem, closely related to that of AHP, and attempt to convince the reader that it is

E-mail address: ojboard@amherst.edu (O. Board).

coherent. In particular, we show within the context of a formal epistemic model that AHP's concept of action-optimality can be characterized in much the same way as Aumann and Brandenburger's analysis of Nash equilibrium (Aumann and Brandenburger, 1997). In Sections 3 and 4 we discuss two variants of the original problem in the context of the same interpretation. In these variants, we relax the assumption that a player's information sets partition his set of decision nodes. Once this relaxation has been accepted, we show that planning optimality no longer guarantees action-optimality (Section 3), and a driver with more information may do worse than a driver with less (Section 4).

As well as being interesting in their own right, it is hoped that the examples discussed in this paper may shed some light on the analysis of decision problems and games with absent-mindedness. Section 5 suggests a way of using the insights gained to construct a general theory.

2. An epistemic analysis of *The Absent-Minded Driver*

The Absent-Minded Driver describes the attempted journey home of an individual late one night from a bar. There are two exits on the highway and he has to take the second to get home (payoff 4). If he takes the first, he reaches a bad neighborhood (payoff 0), and if he fails to take either he has to stay in a motel at the end of the highway (payoff 1). Unfortunately he cannot distinguish between the two intersections: at the second he cannot remember passing the first, and at the first he thinks he may be at the second but have forgotten passing the first. The extensive form of this decision problem is shown below, with the first and second intersections labelled 1 and 2, respectively (Fig. 1).

AHP distinguish between a *planning* stage and an action stage. At the planning stage, the driver is sitting in the bar, wondering what to do when he is on the road. Assuming randomization is allowed, his aim is to maximize $(1-x) \cdot 0 + x(1-x) \cdot 4 + x^2$, where x is the probability of CONT at each intersection. The (unique) planning-optimal strategy is to set $x = \frac{2}{3}$, i.e. 'CONT with probability $\frac{2}{3}$, EXIT with probability $\frac{1}{3}$ '.

The action stage considers the driver's decision when he is on the road. AHP make two observations (p. 104):




Fig. 1. The Absent-Minded Driver.

- First, the driver makes a decision at *each* intersection through which he passes. Moreover, when at one intersection, he can determine the action *only there*, and *not* at the other intersection—where he is not.
- Second, since he is in completely indistinguishable situations at the two intersections, whatever reasoning obtains at one must also obtain at the other, and he is aware of this.

Letting x and y be the probabilities of CONT at the current intersection and at the other intersection, respectively, they show that the expected utility at the action stage is given by

$$h(x, y) = \frac{1}{1+y} [(1-x) \cdot 0 + x(1-y) \cdot 4 + xy \cdot 1] + \frac{y}{1+y} [(1-x) \cdot 4 + x \cdot 1],$$

where $1/(1+y)$ is the probability of being at the first intersection conditional on the fact that he is at *an* intersection. The driver maximizes $h(x, y)$ with respect to x holding y fixed (first observation); the value of x chosen in this way must turn out equal to y (second observation). A strategy x^* is described as *action-optimal* if it satisfies these two conditions, i.e. if it maximizes $h(x, x^*)$.

AHP observe that action-optimality looks very much like an equilibrium concept: formally, (x^*, x^*) “is a symmetric Nash equilibrium in the (symmetric) game between ‘the driver at the current intersection’ and ‘the driver at the other intersection’ (the strategic form of the game with payoff functions h)”. It is interesting and will prove instructive to see how far we can take this analogy, and in particular whether we can provide epistemic conditions for action-optimality that correspond to the epistemic conditions for Nash equilibrium offered by [Aumann and Brandenburger \(1997\)](#). Such analysis has proved successful in elucidating a range of game-theoretic solution concepts, and may help us to achieve a better understanding of action-optimality.

To construct an epistemic model of *The Absent-Minded Driver*, we must amend the standard model¹ in two ways. The first problem to be addressed is that a *state* describes each play of the game in terms of strategies. Usually, each state specifies the pure strategy choice of each player: it is assumed that players do not explicitly randomize. Rather, “...each player chooses some definite action. But other players need not know which one, and the mixture represents their uncertainty, their conjecture about his choice” ([Aumann and Brandenburger, 1997; p. 1162](#)). Our driver moves at only one information set, where two actions are available; so he has only two strategy choices available to him: CONT and EXIT. But neither of these strategies allows him to reach home! So we have a feasible outcome which does not occur at any state in the model: the implication is that the driver thinks it is *impossible* for him to make it home. Allowing explicit randomization, however, with every state specifying a behavior strategy for each player, is not a satisfactory resolution. This generates models in which CONT is chosen with the same probability at both intersections in every state. But AHP’s analysis of the action stage asked the driver to choose a probability of CONT at the current intersection, taking his

¹ The term *standard model* is perhaps not well-defined in this context, since a number of variants are used in the literature. But the comments to follow apply generally.

(probabilistic) behavior at the other intersection as fixed. Action optimality requires that his choices at the two intersections are the same, but this is a contingent fact and not a necessary truth.

In the standard analysis of extensive form games, information sets play two roles: they tell us what the players know whenever they are on move, and they are used to define strategies, the objects of choice for the players. These two roles are logically distinct, and indeed AHP's action stage rejects the second. The driver chooses between the actions available to him at the current intersection only, relaxing the constraint that $x = y$. To construct an epistemic model of the action stage, therefore, we allow the states to specify different actions at each intersection, "because there are two independent decisions" (p. 113). As we shall see, this step allows an interpretation of action-optimality in which explicit randomization is not required (though this is *not* the interpretation of AHP, who state explicitly that randomized behavior is necessary at the action stage).

The second problem with the standard epistemic analysis arises because the states make no explicit reference to time. This is unproblematic in games with no absent-mindedness because, at any given information set, uncertainty about the physical world relates only to which strategies have been, are being, and will be played at the other information sets. In *The Absent-Minded Driver*, however, if the driver chooses CONT at the first intersection then the same information set is reached twice. Indeed, this feature is the defining property of absent-mindedness. Even if he is certain which actions he takes at each intersection, he may still be unsure which intersection he is at. To reflect this uncertainty we need a state corresponding to each intersection, even if the same action is chosen at both.

The most obvious extension of the model, which simply adds a specification of time to each state, does not provide an ideal resolution. For we would like our model to reflect the passage of time, and to describe the player's beliefs at both stages. As we have explained, in the standard model the states are timeless; each state can, therefore, describe beliefs at every information set that is reached. But this is no longer possible once time is introduced, since a given state describes one time only. The *systems* of Halpern and Fagin (1989) provide a natural remedy to this problem.² A system consists of a collection of *runs*; each run is associated with an action profile, which describes what action is taken or would be taken at every decision node of the game. The action profile performs two functions: it describes the way the game is actually played, and it provides a set of counterfactuals for evaluating the payoffs if the action taken at any node deviated from the specified action. These counterfactuals are essential for determining the rationality of each action. Each run is divided into a set of *states*, one for each decision node that is reached along the specified play path. So at a given state, what happened before and after that state was reached is described by the preceding and succeeding states in the same run (if any).

² Halpern (1997) also applies this model to the analysis of decision problems with absent-mindedness. But his approach is very different from that of the current paper, and he examines a different set of issues. Specifically, his aim is to provide a formal comparison of the different solution concepts that have been proposed, while our aim is to consider what lessons can be learned from several variants of the original problem.




Fig. 2. A system for the Absent-Minded Driver.

A natural system for *The Absent-Minded Driver* has a run corresponding to each feasible action profile. There are, therefore, four runs, which we label according to the corresponding action profile (see above). Notice that although the action profiles $(EXIT_1, EXIT_2)$ and $(EXIT_1, CONT_2)$ specify the same path through the game, we need separate runs for each to represent the potential uncertainty in the mind of the driver, when at the first intersection, about what he would do were the second intersection to be reached. There are six states in total, which we label with the letters a to f . Each state is associated with an action profile (given by the corresponding run) and a time period: at state a , for example, the driver plays EXIT at the first intersection, and would play EXIT if the second intersection were reached; at state e , the driver has already played CONT at the first intersection, and is about to play EXIT at the second.

The driver's information is given by a partition over the states, which reflects the information structure of the game, along with the additional assumption that he knows his current action. So states a , b and e , where he chooses EXIT, are contained in one cell, and states c , d and f , where he chooses CONT, are contained in the other. Events are sets of states, as usual. For example, the event that he will make it home is the set $\{c, e\}$, and the event that he is currently at the first intersection is the set $\{a, b, c, d\}$ (Fig. 2).

It will be useful to distinguish between what the driver knows when he is at the first intersection (agent 1) and when he is at the second (agent 2): we shall be interested in what agent 1 knows about agent 2 and vice versa. Let R denote the set of runs, and W the set of states. A run is an ordered tuple of states, so a state can be identified by the run to which it belongs and the time at which it occurs within that run (for example, $(r, 1)$ denotes the first state in run r). Agent i knows that E at run r if he gets to move and E holds at every state he considers possible when the true state is (r, i) . Formally,

$$K_i(E) = \{r \in R \mid (r, i) \in W \text{ and } \mathcal{P}(r, i) \subseteq E\},$$

where $\mathcal{P}(r, i)$ is the cell of the driver's information partition containing state (r, i) .

The probabilistic beliefs of each agent are derived from a prior probability measure P_i on R , the set of runs.³ P_i can be thought of as representing the hypothetical ex ante beliefs of agent i about which action profile will be chosen. When on move, of course, the agent may be unsure also about which intersection has been reached, i.e. about the state as well as the run. A measure over the states, p_i , is derived from P_i by assigning to each state the measure of the run containing it,⁴ and normalizing. The probabilistic beliefs of agent i at state w (when defined) are then given by

$$p_i^w(E) = \frac{p_i(\mathcal{P}(w) \cap E)}{p_i(\mathcal{P}(w))}.$$

We are now in a position to use our system to provide an epistemic analysis of the action-optimal strategy ‘CONT with probability $\frac{2}{3}$ ’. Let P be the prior on R induced by this behavior strategy, so that $P(\text{EXIT}_1, \text{EXIT}_2) = \frac{1}{9}$, $P(\text{EXIT}_1, \text{CONT}_2) = \frac{2}{9}$, $P(\text{CONT}_1, \text{EXIT}_2) = \frac{2}{9}$, and $P(\text{CONT}_1, \text{CONT}_2) = \frac{4}{9}$. The measure p over the states is derived from P as described above: $p(a) = 1/15$, $p(b) = p(c) = p(e) = 2/15$, and $p(d) = p(f) = 4/15$. Finally, let $p_1 = p_2 = p$. It is easy to see that every agent is rational at every run. Consider, for example, agent 2 at state e . Letting $[t = i]$ denote the states at which agent i is on move (i.e. the states where the driver is at the i th intersection), we have

$$p_2^e([t = 1]) = \frac{p_2(\{a, b, e\} \cap \{a, b, c, d\})}{p_2(\{a, b, e\})} = \frac{\frac{1}{15} + \frac{2}{15}}{\frac{1}{15} + \frac{2}{15} + \frac{2}{15}} = \frac{3}{5},$$

$$p_2^e([t = 2]) = \frac{p_2(\{a, b, e\} \cap \{e, f\})}{p_2(\{a, b, e\})} = \frac{\frac{2}{15}}{\frac{1}{15} + \frac{2}{15} + \frac{2}{15}} = \frac{2}{5}.$$

These are precisely the probabilities assigned to the first and second intersections by AHP’s action-optimal driver, and leave the driver indifferent between CONT and EXIT: it is certainly rational for our agent 2 to choose EXIT.

And since every agent is rational at every run, every agent knows at every run that the agent on move at the other intersection (if there is one) is rational. These two conditions correspond to the first two requirements of [Aumann and Brandenburger’s \(1997\)](#) epistemic characterization of mixed-strategy Nash equilibrium. But Aumann and Brandenburger also require that the players know each other’s conjectures (indeed it is these conjectures that are interpreted as the mixing probabilities in the mixed-strategy Nash equilibrium). The immediate analog of this condition is not satisfied here. Consider again agent 2 at state e . What does he know about agent 1’s conjecture about CONT_2 ? Agent 2 is at state e , so he considers states a , b and e to be possible. At states a and b , we have $p_1^a(\text{CONT}_2) = p_1^b(\text{CONT}_2) = \frac{2}{5}$, while if state e is the true state, agent 1’s beliefs when on move earlier in the run are given by $p_1^c(\text{CONT}_2) = \frac{4}{5}$. The reason for this

³ It may seem strange to assign different priors to two agents, who after all represent the same person. In the context of the current example we agree, since the structure of the game is such that both agents possess the same information. But we see or reason why two agents who move at different information sets should possess the same prior, unless one wishes to adopt the common prior assumption in general. Systems do not model the development of the player’s beliefs over time as the game progresses, but merely describe those beliefs: the information partition of agent 2 is not derived from the partition of agent 1, though of course the two may coincide.

⁴ If a given run is realized then each state in that run is certain to occur.

anomaly is that agent 2 does not know who he is! Although he knows his own beliefs about the likelihood of CONT_2 , he does not know whether he is agent 1 or agent 2, and the two agents' beliefs differ.

AHP are aware of this problem, and talk about each agent's beliefs about 'the other intersection'. Agent 2 at state e does not know which agent he is, but we can calculate his beliefs about what the agent at the other intersection is doing if he is agent 1 and if he is agent 2 separately.

From the point of view of agent 2, the agent at the other intersection plays CONT if (a) $t = 1$ and CONT_2 ; or (b) $t = 2$ and CONT_1 . The probability he assigns to this event is given by

$$p_2^e((t = 1] \cap \text{CONT}_2) \cup ((t = 2] \cap \text{CONT}_1)) = \frac{\frac{2}{15} + \frac{2}{15}}{\frac{1}{15} + \frac{2}{15} + \frac{2}{15}} = \frac{4}{5}.$$

Carrying out the analogous calculation at every other state we find that the agent on move always assigns a probability of $\frac{4}{5}$ to the agent at the other intersection playing CONT; thus at every run each agent knows the conjecture of the agent at the other intersection about what the agent at the other intersection will do. This corresponds to Aumann and Brandenburger's requirement that the agents know each other's conjectures, and suggests that it may be possible to interpret action-optimality in the same way as mixed-strategy Nash equilibrium. We do not assume that the agents choose explicitly random strategies: the need for randomization is removed because at the action stage the agents are indifferent between CONT and EXIT. Rather the mixing is to be thought of as representing uncertainty in the mind of the agent at the other intersection. Although there are states in which the two agents choose different actions, AHP's second observation still holds: the agents reason the same way each intersection.

Before moving on, we address one final issue. The action-optimal strategy calculated by AHP is to play CONT with probability $\frac{2}{3}$; and yet we have just calculated that each agent believes that the other plays CONT with probability $\frac{4}{5}$. But this difference does not reflect time inconsistency of the form discussed by [Piccione and Rubinstein \(1997\)](#). Rather the two figures represent the agents' *ex ante* and *interim* beliefs. The fact that an agent gets a chance to move gives him information about what might have happened at the other intersection. In particular, he can rule out the situation in which he is at the second intersection and the agent at the first intersection played EXIT: in this case the second intersection would not have been reached. $\frac{4}{5}$ is the probability that the agent at the other intersection plays CONT conditional on the current agent reaching an intersection.

3. Tiredness and non-partitional information sets

Suppose that the driver is not absent-minded but merely tired. At the first intersection he is still sufficiently alert to be certain he really is at the first intersection, but by the time he reaches the second he cannot recall whether he has passed the first or not. His information sets at the two intersections overlap but do not coincide: using the obvious

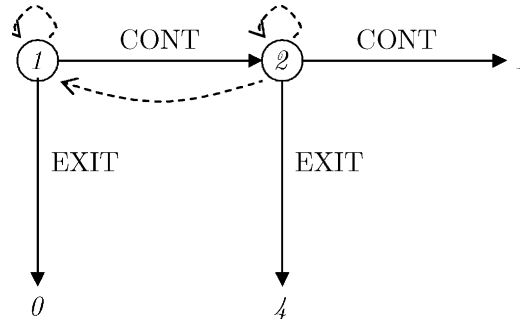


Fig. 3. The tired driver.

notation, we have $I_1 = \{1\}$ and $I_2 = \{1, 2\}$. These information sets do not partition the set of decision nodes, and so cannot be represented in the standard way.⁵ Rather, we use dotted arrows to show which intersections each agent considers possible when on move (Fig. 3).

It may seem that the agent on move at intersection 2 should be able to figure out where he is: if he were at intersection 1 he would know he was there; and if he knows the information structure of the game, he knows this. So since he is unable to tell whether he is at intersection 1 or 2, he should infer that he must be at intersection 2. But this argument assumes that he knows that he does not know whether he is at intersection 1 or 2, while the information sets tell us only that he does not know. Implicitly we are assuming that the agent does not satisfy negative introspection: he does not know everything he does not know.⁶

The Tired Driver has a unique planning-optimal strategy. A strategy specifies an action for each information set at which the player is on move, and since the information sets at the two intersections are distinct, we can specify a different action at each and guarantee that the driver reaches home. His planning-optimal strategy, then, is to play CONT at I_1 and EXIT at I_2 . To investigate action-optimality, we construct a system for the game, again assuming that there is one run corresponding to each action profile (Fig. 4).

The dotted arrows work in the same way as the arrows in the game itself and tell us, for each state, which states the agent on move at that state considers possible. If two nodes are enclosed in a box then each is considered possible from the other. Thus $\mathcal{P}(b) = \{a, b\}$, and $\mathcal{P}(e) = \{a, b, e\}$, for example. The time arrows are omitted to avoid clutter.

The planning-optimal strategy CONT at I_1 and EXIT at I_2 induces the following prior on R : $P(\text{EXIT}_1, \text{EXIT}_2) = P(\text{EXIT}_1, \text{CONT}_2) = P(\text{CONT}_1, \text{CONT}_2) = 0$; $P(\text{CONT}_1, \text{EXIT}_2) = 1$. According to this prior, only run $(\text{CONT}_1, \text{EXIT}_2)$ is reached with positive probability, and at this run (a) both agents are rational; (b) they both

⁵ Games with non-partitional information structures are discussed in more detail in Board (2002), where it is shown how they can be used to model deception.

⁶ In fact, it follows from a well-known result of modal logic that the agent satisfies positive introspection if and only if the dotted arrows which represent the agent's information sets describe a transitive binary relation on the decision nodes (as they do here); and he satisfies negative introspection if and only if they describe a Euclidean binary relation (i.e. if $x \rightarrow y$ and $x \rightarrow z$ then $y \rightarrow z$).

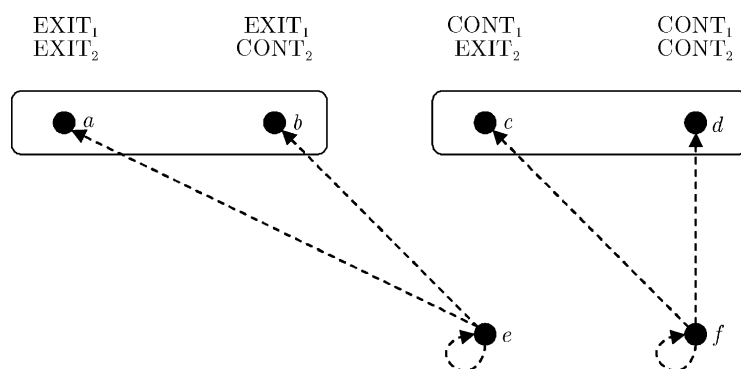


Fig. 4. A system for the tired driver.

know the agent at the other intersection is rational; and (c) they both know the conjecture of the agent at the other intersection about what the agent at the other intersection will do. But there is something strange about agent 2's beliefs at state *e*. Agent 2 assigns $p_2^e([t = 1]) = 0$ and $p_2^e([t = 2]) = 1$, so that he believes with probability 1 that he is at intersection 2. The reason is that he knows that agent 1 plays CONT and agent 2 plays EXIT; and since he is playing EXIT, he must be agent 2. But by the same token, presumably if he had decided to play CONT, he would have believed he was agent 1! Even if we do not wish to take literally the idea that the agents' priors represent their beliefs at some pre-decision reasoning stage, it is important to provide a coherent story of how their posteriors are derived. A story which allows an agent to draw inferences from his own action choice about which intersection has been reached is not coherent. It is this kind of reasoning which leads evidentialist decision theorists to take one box in Newcomb's problem.

Rather, AHP show that the probabilities an agent assigns to each intersection should depend only on his belief about what the agent at the other intersection is doing. If we want to take seriously the idea that agent 2 cannot distinguish between intersections 1 and 2, then the same reasoning should apply to him. It follows that his prior must assign the same probability to EXIT₁ as to EXIT₂, and that these probabilities must be independent of each other—only then will his posterior not depend on his actual choice.⁷ So we must assign different priors to agent 1 and agent 2. The following priors satisfy the Aumann and Brandenburger conditions and the additional AHP independence condition:

$$\begin{aligned}
 P_1(\text{EXIT}_1, \text{EXIT}_2) &= P_1(\text{EXIT}_1, \text{CONT}_2) = 0; & P_1(\text{CONT}_1, \text{EXIT}_2) &= \frac{1}{3}; & P_1(\text{CONT}_1, \text{CONT}_2) &= \frac{2}{3} \\
 P_2(\text{EXIT}_1, \text{EXIT}_2) &= \frac{1}{9}; & P_2(\text{EXIT}_1, \text{CONT}_2) &= P_2(\text{CONT}_1, \text{EXIT}_2) = \frac{2}{9}; & P_2(\text{CONT}_1, \text{CONT}_2) &= \frac{4}{9}
 \end{aligned}$$

⁷ This also implies that agent 2 cannot recall what he planned to do at the ex ante stage; for as we have seen, the planning optimal strategy involves playing CONT at intersection 1. A referee has commented that this is a rather strong assumption. But it seems appropriate in a context in which the driver cannot recall what he actually did (if anything). Further, even if the driver can recall what he planned to do, he may not be sure that he succeeded in carrying out this plan.

The action-optimal strategy is calculated by taking the prior of the agent on move at each intersection over actions at that intersection:

I_1 : CONT;
 I_2 : CONT with probability $\frac{2}{3}$; EXIT with probability $\frac{1}{3}$.

Before moving on, we make a few observations about this result. First, agent 2's strategy is exactly the same as it was in *The Absent-Minded Driver*. The reason is that he has exactly the same information about which intersection has been reached (i.e. no information); although agent 2 in *The Tired Driver* knows that agent 1 is better informed, he cannot use this information because he does not know his own beliefs. Second, the planning-optimal strategy in *The Tired Driver* is not action-optimal. This provides a counter-example to Proposition 3 of Piccione and Rubinstein (1997), and shows that their result relies on partitional information sets. Finally, the action-optimal strategy in *The Tired Driver* yields higher ex ante expected payoff ($\frac{1}{3} \cdot 4 + \frac{2}{3} \cdot 1 = 2$) than the action-optimality strategy in *The Absent-Minded Driver* ($\frac{1}{3} \cdot 0 + \frac{2}{9} \cdot 4 + \frac{4}{9} \cdot 1 = 1\frac{1}{3}$). The tired driver benefits from his attentiveness at the first intersection by avoiding the worst payoff of 0. But we shall see in Section 4 that more information is not always better than less.

4. More is not always better

Consider the following decision problem, where the driver lives off the third intersection on the road (Fig. 5).

We shall examine two different information structures:

$A : I_1 = \{1\}; I_2 = \{2\}; I_3 = \{2, 3\};$
 $B : I_1 = \{1\}; I_2 = \{2\}; I_3 = \{1, 2, 3\}.$

Driver A is better informed than driver B in a strict sense: driver A is able (correctly) to exclude some intersections that driver B cannot, but the reverse is not true. It is clear that the planning-optimal strategy in both versions of the game is the same, and achieves the maximum payoff since the driver has a different information set, and therefore can choose a different action, at every intersection: CONT at I_1 , CONT at I_2 , EXIT at I_3 .




Fig. 5. A longer road.

But in neither case is this strategy action-optimal. Look first at driver A. Agents 1 and 2 are perfectly informed, and so agent 1 will play CONT as long as the expected payoff from doing so is greater than 2; and agent 2 will always play CONT, since it is a dominant strategy for him. If he gets a chance to move, agent 3 faces precisely the same decision problem as agent 2 in *The Tired Driver*, and his unique action-optimal strategy is to play CONT with probability $\frac{2}{3}$. This leaves agent 1 indifferent between playing CONT and playing EXIT: his expected payoff either way is $\frac{1}{3} \cdot 4 + \frac{2}{3} \cdot 1 = 2$. So we have a continuum of action-optimal strategies (CONT with probability x ($0 \leq x \leq 1$) at I_1 ; CONT at I_2 ; CONT with probability $\frac{2}{3}$ at I_3), all yielding the same ex ante expected payoff of 2.

Agents 1 and 2 of driver B are also perfectly informed; as before agent 1 will play CONT as long as the expected payoff from doing so is greater than 2, and agent 3 will always play CONT. So suppose that intersection 3 is reached. We argued in Section 3 that an agent who is unable to distinguish between two intersections must have identical and independent prior probability distributions over the actions available at each, and the same applies to an agent who is unable to distinguish between three intersections. His prior must, therefore, be of the form:

$$P_3(\text{EXIT}_1, \text{EXIT}_2, \text{EXIT}_3) = (1 - x)^3;$$

$$P_3(\text{EXIT}_1, \text{CONT}_2, \text{EXIT}_3) = x(1 - x)^2;$$

$$P_3(\text{EXIT}_1, \text{EXIT}_2, \text{CONT}_3) = x(1 - x)^2;$$

$$P_3(\text{EXIT}_1, \text{CONT}_2, \text{CONT}_3) = x^2(1 - x);$$

$$P_3(\text{CONT}_1, \text{EXIT}_2, \text{EXIT}_3) = x(1 - x)^2;$$

$$P_3(\text{CONT}_1, \text{CONT}_2, \text{EXIT}_3) = x^2(1 - x);$$

$$P_3(\text{CONT}_1, \text{EXIT}_2, \text{CONT}_3) = x^2(1 - x);$$

$$P_3(\text{CONT}_1, \text{CONT}_2, \text{CONT}_3) = x^3;$$

where x is the prior probability he assigns to CONT at each intersection. It is easy to see that $x = 0$ is an action-optimal strategy. For if $x = 0$, when on move he believes with probability 1 that he is at the first intersection,⁸ and he believes with probability 1 that the agent at intersection 2 will play EXIT: he should, therefore, play EXIT. And in fact $x = 0$ is the only action-optimal strategy for agent 3. $x = 1$ is not action-optimal, since if $x = 1$ agent 3 must assign equal probability to intersections 1, 2 and 3, and he should, therefore, play EXIT. Nor is there an action-optimal strategy $0 < x < 1$. For such a strategy to be action-optimal, it must be that the *interim*⁹ expected payoff from playing CONT is the same

⁸ Note that this belief is not derived from beliefs about his *own* action choice, but rather from his beliefs about what the agents at the other intersections are doing. Hence we are not committing the evidentialist fallacy discussed in Section 3.

⁹ i.e. the expected payoff of agent 3 according to the beliefs he has when on move.

as the *interim* expected payoff from playing EXIT:

$$\begin{aligned} & \frac{1}{1+x+x^2}[x(1-x)\cdot 4+x^2\cdot 1] + \frac{x}{1+x+x^2}[(1-x)\cdot 4+x\cdot 1] + \frac{x^2}{1+x+x^2}\cdot 1 \\ &= \frac{1}{1+x+x^2}\cdot 2 + \frac{x}{1+x+x^2}\cdot 0 + \frac{x^2}{1+x+x^2}\cdot 4 \Rightarrow 9x^2 - 8x + 2 = 0 \end{aligned}$$

where $1/(1+x+x^2)$, $x/(1+x+x^2)$ and $x^2/(1+x+x^2)$ are the posterior probabilities agent 3 assigns to intersections 1, 2 and 3. But this equation has no real roots, and therefore there are no values of x which leave agent 3 indifferent between playing CONT and EXIT. So we have found an action-optimal strategy for driver *B*: CONT at I_1 ; CONT at I_2 ; EXIT at I_3 . This yields an ex ante expected payoff of 4, twice as much as driver *A*!

5. Conclusions

The contents of this paper do not provide a general theory of equilibrium in games with absent-mindedness, or even of action-optimality in decision problems with absent-mindedness. Such a theory would offer a formal, game-theoretic definition of action-optimality, and show that it could be characterized by a particular set of epistemic conditions in the sense that (a) whenever those conditions were satisfied, the players would play an action-optimal strategy profile, and (b) whenever the players were to play an action-optimal strategy profile, they could be modeled as if satisfying those conditions (by an appropriate system of the game). This is a task for future research. The current aim has been to shed some light on what this theory might look like, and to show by means of simple examples that action-optimality has some interesting properties.

Acknowledgements

I would like to thank Giacomo Bonanno, Bob Stalnaker and an anonymous referee, whose comments lead to substantial improvements in the exposition of this paper.

References

- Aumann, R.J., Brandenburger, A., 1997. Epistemic conditions for Nash equilibrium. *Econometrica* 63, 1161–1180.
- Aumann, R.J., Hart, S., Perry, M., 1997. The absent-minded driver. *Games and Economic Behavior* 20, 102–116.
- Board, O.J., 2002. The deception of the Greeks: generalizing the information structure of extensive form games. *Greek Economic Review* 22, 1–16.
- Halpern, J.Y., 1997. On ambiguities in the interpretation of game trees. *Games and Economic Behavior* 20, 66–96.
- Halpern, J.Y., Fagin, R., 1989. Modelling knowledge and action in distributed systems. *Distributed Computing* 3, 159–179.
- Piccione, M., Rubinstein, A., 1997. On the interpretation of decision problems with imperfect recall. *Games and Economic Behavior* 20, 3–24.