

A Note On Imperfect Recall

by Ken Binmore

Economics Department

University College London

Gower Street, London WC1E 6BT

England

Abstract. The Paradox of the Absent-Minded Driver is used in the literature to draw attention to the inadequacy of Savage's theory of subjective probability when its underlying epistemological assumptions fail to be satisfied. This note suggests that the paradox is less telling when the uncertainties involved admit an objective interpretation as frequencies.

A Note On Imperfect Recall¹

by Ken Binmore
Economics Department
University College London
Gower Street, London WC1E 6BT
England

Ce que j'ai appris, je ne le sais plus. Le
peu que je sais encore, je l'ai deviné.

Chamfort, *Maximes et Pensées*, 1795

1 Introduction

Von Neumann and Morgenstern [17] proposed modeling bridge as a two-person, zero-sum game in which each partnership is one of the two players. Modeled in this way, bridge becomes a game of imperfect recall, because the players forget things they knew in the past. For example, when bidding as North, the player representing the North-South partnership must forget the hand he held when bidding as South. A bank choosing a decentralized lending policy for its hundred branches confronts a similar problem. The boss can imagine himself behind each manager's desk as he interviews clients, but then he must forget the earlier lending decisions he made while sitting behind other managers' desks. The decision problem he faces is therefore one of imperfect recall.

Since Kuhn [9] pointed out that mixed and behavioral strategies are interchangeable only in games of perfect recall, little attention has been paid to the difficulties that can arise when recall is imperfect.² The orthodox approach has been to regard a person with imperfect recall as a team of agents who have identical preferences but different information. In the style of Selten [15], each agent is then treated as a distinct player in the game used to model the problem. For example, the orthodox approach models bridge as a four-player game with two teams and the banking problem as a hundred-player game with one team.

As Gilboa [6] confirms, the consensus in favor of the team approach is very strong. Nevertheless, a recent paper of Piccione and Rubinstein [12] has revived interest in the problem of decision-making with imperfect recall. Their emphasis in this paper is on the straightforward psychological fact that most people know

¹Support from the Economic and Social Research Council under their "Beliefs and Behaviour" Programme L 122251024 is gratefully acknowledged.

²Notable exceptions are Isbell [7] and Alpern [1].

that they will forget things like telephone numbers from time to time. The orthodox approach dismisses such folk as irrational and thereby escapes the need to offer them advice on how to cope with their predicament. But I agree with Piccione and Rubinstein that people who know themselves to be absent-minded can still aspire to behave rationally in spite of their affliction. However, my own interest in the imperfect recall problem is mostly fuelled by the difficulties with the team approach outlined in Binmore [4].

Players stay on the equilibrium path of a game because of their beliefs about what would happen if they were to deviate. But how do players know what would happen after a deviation? The orthodox approach treats a player as though he were a team of agents, one for each information set at which he might have to decide what action to take. Such a viewpoint de-emphasizes the inferences that a player's opponents are likely to make about his thinking processes after he deviates. One agent in a team may have made a mistake, but why should that lead us to think that other agents in the same team are liable to make similar mistakes? Traditionalists see no reason at all, and hence their allegiance to backward induction and similar solution concepts. However, a theory that treats a player as a team of independently acting agents is unlikely to have any realistic application, because we all know that real people are liable to repeat their mistakes. If I deviate from the equilibrium path, it would therefore be stupid for my opponents not to make proper allowance for the possibility that I might deviate similarly in the future.³ The team approach is therefore not without its difficulties even in games of perfect recall. It therefore seems an unlikely panacea for imperfect recall problems.

Piccione and Rubinstein [12] do not claim to provide a theory of rational decision-making under imperfect recall. They seek only to comment on some of the issues that would need to be addressed in formulating such a theory. I am even less ambitious in that I shall simply be commenting on some of their comments. The problem of time-inconsistency raised by the one-player game that they aptly describe as the Paradox of the Absent-Minded Driver is particularly interesting.

2 Paradox of the Absent-Minded Driver

The general issue of imperfect recall in games is discussed in my *Fun and Games* (Binmore [3]). Such textbooks explain how to interpret representations of imperfect recall problems like that shown in Figure 1(a).⁴ One may imagine an

³Binmore [4] argues that one needs an explicit algorithmic model of the reasoning Processes of a player in order to take account of such considerations. Finite automata have been used for this purpose in a number of papers. However, as Rubinstein [13] notes, one cannot model players as finite automata without introducing imperfect recall problems.

⁴Von Neumann and Morgenstern [17] excluded cases like this by requiring that no play of a game should pass through an information set more than once. However, it has now become customary to accept such cases as particularly challenging examples of games of imperfect

absent-minded driver who must take the second turning on the right if he is to get home safely for a payoff of 1. If he misses his way, he will find himself in an unsafe neighborhood and receive a payoff of 0.

The driver's difficulty lies in the fact that he is absent-minded. At each exit, he forgets altogether what has happened previously in his journey. Since both exits look entirely the same, he is therefore unable to distinguish between them.⁵

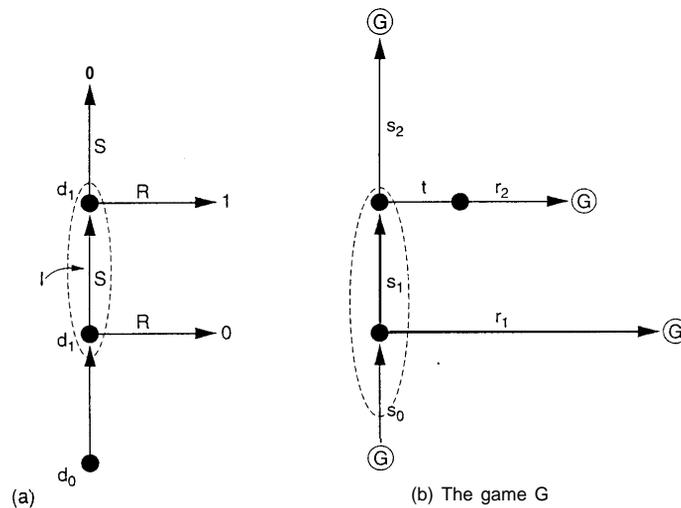


Figure 1: The absent-minded driver's problem

The driver has two pure strategies for this one-player game of imperfect recall, R and S . The use of either results in a payoff of 0. It follows that the same is true of any mixed strategy. However, in such games of imperfect recall, one can achieve more by using a behavioral strategy. A behavioral strategy requires the driver to mix between R and S each time he finds himself called upon to make a decision.⁶ Let $b(p)$ be the behavioral strategy in which R is chosen with probability $1 - p$ and S is chosen with probability p . A driver

recall.

⁵One can, of course, invent ways in which he could supplement his memory. For example, he might turn on the radio on reaching an exit. He could then distinguish between the exits by noting whether his radio is on or off. But the introduction of such expedients is not allowed.

⁶By contrast, a mixed strategy requires a player to randomize over his Pure strategies once and for all before the game is played.

who uses $b(p)$ obtains an expected payoff of $p(1 - p)$, which is maximized when $p = \frac{1}{2}$. According to this analysis, his optimal behavioral strategy is therefore $b(\frac{1}{2})$, which results in his receiving an expected payoff of $\frac{1}{4}$.

The paradox proposed by Piccione and Rubinstein [12] hinges on the time at which the driver chooses his strategy. The argument given above takes for granted that the driver chooses $b(\frac{1}{2})$ before reaching node d_1 , and that he can commit himself not to revise this strategy at a later date. But such an attitude to commitment is not consistent with contemporary thinking in game theory. In particular, Selten's [15] notion of a perfect equilibrium, together with all its successors in the literature on refinements of Nash equilibrium, assumes that players will always be re-assessing their strategy throughout the game.⁷ More precisely, the orthodox view is not that players *cannot* make commitments, but, if they can, their commitment opportunities should be modeled as formal moves in the game they are playing. However, once their commitment opportunities have been incorporated into the rules of the game, then the resulting game should be analyzed without attributing further commitment powers to the players.

So what happens in the absent-minded driver's paradox if the driver is not assumed to be committed to $b(-)$? Following Piccione and Rubinstein [12], let us assume that he reaches the information set I and remembers that he previously made a plan to choose $b(-)$ on reaching I . He then asks himself whether he wants to endorse this strategy now that he knows he has reached the information set I and hence may either be at d_1 or d_2 . If he attaches probability $1 - q$ to the event of being at d_1 and probability q to the event of being at d_2 , then choosing $b(p)$ at I results in a payoff

$$\pi = (1 - q)p(1 - p) + q(1 - p) \quad (1)$$

This payoff is maximized when $p = (1 - 2q)/2(1 - q)$. The driver will therefore only choose $p = \frac{1}{2}$ at I if he believes that $q = 0$. That is to say, in order that a time-inconsistency problem not arise, it is necessary that the driver deny the possibility of ever reaching the second exit. But to deny the possibility of reaching the second exit is to deny the possibility that he can ever get home!

3 Whence q ?

To make progress with the Paradox of the Absent-Minded Driver, it is necessary to ask how the driver came to believe that the probability of being at the second exit d_2 is q . This question forces us in turn to face a philosophical question about the nature of probability. In the terminology of Binmore [5, p.265], is the driver's probability theory logistic, subjective or objective? A logistic theory

⁷McClennen [10] is one of a number of philosophers who insist that rationality includes the facility to commit oneself to perform actions in the future under certain contingencies that one's future self would regard as suboptimal.

treats a probability as the rational degree of belief in an event justified by the evidence. The subjective theory of Savage [14] is the basis for the familiar Bayesian orthodoxy of economics.⁸ An objective theory regards the probability of an event as its long-run frequency.

The most satisfactory interpretation for q in the Paradox of the Absent-Minded Driver would be logistic, in the style attempted by Keynes [8]. However, I think it uncontroversial that no theory of this type has yet come near being adequate. My guess is that most economists would take for granted that q is to be interpreted as a subjective probability à la Savage [14]. But there are major difficulties in such an interpretation. In the first place, Bayesian epistemology—as described in Chapter 10 of Binmore [3]—fails in the absent-minded driver's problem. ⁹Secondly, if the postulates of Savage's theory are to make sense, it is vital that the action space A , the space B of states of the world, and the space C of consequences have no relevant linkages other than those incorporated explicitly in the function $f : A \times B \rightarrow C$ that determines how actions and states together determine consequences (Binmore [5, p.310]). But it is of the essence in the Paradox of the Absent-Minded Driver that states are not determined independently of actions. A rational driver's beliefs about whether he is at the first exit or the second must surely take account of his current thinking about the probability at which he would turn right if he were to reach an exit.

Personally, I think that the most important role for paradoxes like that of the absent-minded driver is to focus attention on the inadequacies of our current logistic and subjective theories of probability. However, I have nothing particularly original to propose on either front, and so follow Piccione and Rubinstein [12] in this paper by turning to the interpretation of q as an objective probability y .

If q is to be interpreted objectively, one must imagine that the driver faces the same problem every night on his way home from work. After long enough, it is then reasonable to regard the ratio of the number of times he arrives at the second exit to the total number of times he arrives at either exit as a good approximation to the probability q . Of course, this frequency will be determined by how the driver behaves when he reaches an exit. If the driver always continues straight on with probability P at an exit, then the number of times he reaches

⁸Notice that I do not identify Bayesianism with Savage's theory. I distinguish those who subscribe to Savage's view from followers of Bayesianism by calling the former Bayesians and the latter Bayesianismists (Binmore [2,5]). Bayesianismists argue that rationality somehow endows individuals with a prior probability distribution, to which new evidence is assimilated simply by updating the prior according to Bayes' Rule. Such an attitude reinterprets the subjective probabilities of Savage's theory as logistic. This may sometimes be reasonable in a small-world context, but Savage [14] condemns such a procedure as "ridiculous" or "preposterous" in a large-world context.

⁹The epistemology taken for granted by Bayesian decision theory is simply that a Person's knowledge can be specified by an information partition that becomes more refined as new data becomes available. Binmore [3, p.457] discusses the absent-minded driver's problem explicitly in this connexion.

d_2 will be a fraction P of the number of times he reaches d_1 . It follows that

$$q = P/(1 + P) \text{ and } 1 - q = 1/(1 + P). \quad (2)$$

One school of thought advocates writing the values for q and $1 - q$ from (2) into (1) and then setting $p = P$ to obtain $\pi = 2p(1 - p)/(1 + p)$. The result is then maximized to yield the optimal value $p = \sqrt{2} - 1$. However, such a derivation neglects the requirement that a decision-maker should maximize expected utility *given* his beliefs. In what follows, I therefore always treat a player's beliefs as fixed when optimizing, leaving only his actions to be determined. However, if (1) is maximized with $q = P/(1 + P)$ held constant, then the maximizing p satisfies $p = (1 - P)/(1 + P)$. A time-inconsistency problem then arises unless $p = P$. Imposing this requirement leads to the equation $p^2 + 2p - 1 = 0$, whose positive solution is $p = \sqrt{2} - 1$ as before. In the next section, I plan to defend this result as the resolution of the Paradox of the Absent-Minded Driver in the case when it is possible to interpret q as a frequency.

4 Repeated Absent-Mindedness

One of the things that game theory has to teach is that difficulties in analyzing a problem can sometimes be overcome by incorporating *all* of the opportunities available to the decision-maker into the formal structure of his decision problem. The need to proceed in this manner has been recognized for a long time in the case of precommitment, and I think it uncontroversial to assert that the orthodox view among game theorists is now that each opportunity a player may have to make a precommitment should be built into the moves of a larger game, which is then analyzed without further commitment powers being attributed to the players.

If the absent-minded driver's decision problem is prefixed with such a move—at which the driver commits himself to choosing a probability p with which to continue straight ahead whenever he reaches an exit—then we have seen that the problem reduces to choosing the largest value of $p(1 - p)$. However, if the problem is presented without a formal commitment move, as in Figure 1(a), then the convention in game theory is to seek an analysis that does not attribute commitment powers to the driver.

It seems to me that the same should go for Piccione and Rubinstein's [12] assumption that the driver is able to remember a decision made in the past about which action he planned to take on encountering an exit. If there are pieces of information that are relevant to the decisions that might be made during the play of a game, then these should be formally modeled as part of the rules of the game. Otherwise, my understanding of the conventions of game theory is that an analysis of the game should proceed on the assumption that the unmodeled information is not available to the players. In particular, we should analyze the absent-minded driver's decision problem as formulated in

Figure 1(a) without assuming that the driver remembers anything at all that is relevant to his decision on arriving at I . One could, of course, introduce a new opening move at which the driver makes a provisional choice of behavioral strategy for the problem that follows. However, personally I think that this issue is something of a red herring. As we all know when we dip our toes in the sea on a cold morning, the plans we made earlier when getting changed are not the plans that determine what we actually do. What we do now is determined by the decision we make now.

If I understand correctly, Piccione and Rubinstein [12] want the driver to remember his original provisional plan so that they have grounds for attributing beliefs to him about whether he is at d_1 or d_2 . But why should his original plan determine his beliefs if he has no reason to believe that he will carry out his original plan once he reaches I ? In my view, this and other issues in the case when q is to be interpreted as a long-run frequency are clarified by explicitly modeling the situation as a *repeated* decision problem, rather than leaving the repetitions to be implicitly understood. One is then led to present the problem as shown in Figure 1(b), where it is to be understood that time recedes into the infinite past as well as into the infinite future. In order to avoid the type of time-inconsistency problems pointed out by Strotz [16], it will be assumed that the driver discounts time according to a fixed discount factor δ ($0 < \delta < 1$).

The label attached to an edge representing an action in Figure 1(b) will also denote the time the driver spends between the nodes joined by the edge. Nothing very much hinges on this point, but I think it helpful to imagine that no time at all is spent at a node, but that the driver does all his thinking while driving between nodes. The information set I in Figure 1(b) is therefore not quite the same as its cousin in Figure 1(a), since it includes not only the nodes d_1 and d_2 , but also the open edge that joins d_0 and d_1 and the open edge that joins d_1 and d_2 .

I envisage the driver moving through the tree reviewing his plan of campaign as he goes. As noted above, whether he actually remembers his previous plan or not seems irrelevant to the question since he will be reiterating all of the considerations each time he reviews his situation. When he reaches an exit, he implements whatever plan he currently has in mind. To keep things simple, I assume that there is a probability of 1 that he will review his plan at least once on each edge of the tree. Notice that, as the problem is set up in Figure 1(b), the driver always remembers what happened when he faced the same problem in the past. Only within I does he forget something, and then the only relevant matter of which he is unaware is his location in I . (He is, of course, not allowed to carry a clock or to employ any other device that would help him to reduce his uncertainty while in I .)

Within this formulation, two simple points seem apparent to me. While suffering dreadfully with a hangover, I can recall swearing never to drink unwisely again. I can also recall making similar resolutions repeatedly in the past. But plans for the future made under such circumstances remain relevant only while

the memory of the hangover is sharp. Once the memory has faded, the joys of convivial company again outweigh the anticipated suffering and one drinks too much again.¹⁰ In brief, if we want to know what someone who is not a regular tippler will do when tempted to overindulge, we need to know how he will view the matter when he is tempted—not immediately after he has overindulged in the past. Similarly, in the problem of the absent-minded driver as modeled in Figure 1(b), it seems to me that only the plan the driver makes while in the set I is relevant.

The second point concerns his beliefs while in the set I . The driver always remembers all his past history up to his most recent entry into the set I . He can therefore compute the frequency P with which he continued straight ahead at exits in the past. This information does not tell him for certain what probabilities he should be using to estimate his location in I now, but it is the strongest possible evidence he could possibly have on this subject.¹¹

Let v be the expected payoff to the driver, given that he is at node d_1 and will always play $b(p)$ at an exit now and in the future. Let w be the expected payoff when he is to use $b(p)$ but is now at node d_2 . Then

$$\begin{aligned} v &= (1 - p)v\delta^{r_1+s_0} + pw\delta^{s_1}, \\ w &= (1 - p)\delta^t + (1 - p)v\delta^{t+r_2+s_0} + pv\delta^{s_2}, \end{aligned}$$

from which it is easy to calculate v and w . One can experiment with various relative values of the time intervals of the model. In some fairly natural cases, it turns out that the driver's informational state is irrelevant, since he would make the same decision at d_1 and d_2 even if he knew his location. The issue is also trivialized by considering the case when the time intervals are fixed and $\delta \rightarrow 1$. However, I plan to consider the limiting case when δ is fixed but $r_1 = r_2 = s_2 = T$ and $T \rightarrow \infty$. This removes the influence of the future repetitions of the problem on the driver's current behavior while still allowing him access to his decisions in the same situations in the past. Under this simplifying hypothesis,

$$\begin{aligned} v &= p(1 - p)v\delta^{s_1+t}, \\ w &= (1 - p)\delta^t. \end{aligned}$$

If the driver knew he were between d and d_1 , he would want to maximize v and so would choose $p = \frac{1}{2}$. If he knew he were between d_1 and d_2 , he would want

¹⁰ Alcoholics Anonymous recommend permanent attendance at group sessions. Is this to keep memories sharp by renewing them vicariously through the experience of others?

¹¹ One might argue that we should model the driver as a computing program. Nothing would then prevent this program from using itself as a subprogram and hence simulating its own thinking processes. Could it not therefore dispense with external information and simply use an introspective analysis to determine what it must have done in the past? Binmore [4] points out that difficulties arise in such cases because of the Halting Problem for Turing Machines. In simple terms, a program that decides what to do on the basis of a prediction of what it is about to do will get into an infinite loop. However, the driver in the formulation given in the text has no need to face the problems that modeling such introspection creates, because he is provided painlessly with the data that the attempt at introspection is intended to generate.

to maximize w and so would choose $p = 0$. But once he is within the set I he does not know his location. But he can easily compute the conditional probability q that the latter case applies to be $q = Ps_1/(s_0 + Ps_1)$. I take $s_0 = s_1 = s$, so that this formula reduces to $q = P/(P + 1)$ as in the analysis of Piccione and Rubinstein [12] discussed in the preceding section. The optimizing calculation for p then proceeds precisely as in their analysis, so that the maximizing p is $(1 - P)/(1 + P)$.

I see no more paradox in the fact that this value of p differs from that the driver would choose outside I than the fact that it differs from the p he would choose if he knew he were approaching d_1 or d_2 . In both cases, he knows less at the time of decision than in the circumstances with which his actual decision is to be compared. In particular, within the set I , he does not know whether he is involved in a decision problem D_1 that starts at d_1 or a decision problem D_2 that starts at d_2 . On the other hand, if any plan he made outside I stood a chance of still being in place when d_1 is reached, then he would be choosing in the knowledge that the problem to be solved is D_1 .

It remains to argue that $p = P$. In accordance with my doubts about backward induction (Binmore [4]), I do not believe that it is possible to tackle this question adequately without considering out-of-equilibrium behavior. However, in the absence of an algorithmic model of the driver's thinking processes, it is only possible to offer a sketch of how a full argument would go.

Begin by considering the possibility that $p \neq P$. The driver now has a problem because the behavior that his calculation recommends for the infinite future is not consistent with his summary of his behavior in the infinite past. He therefore needs some theory of "mistakes" to explain why he did the wrong thing in the past or why he may not actually do what he has calculated to be optimal in the future.¹² When $p \neq P$, the driver will presumably accept that either P is flawed as a prediction of p or $P/(1 + P)$ is flawed as a prediction of q . He will then need to employ more complicated functions of his history to generate the estimates he needs to make a sensible decision.¹³ If these functions move the driver's estimates towards consistent values over time, then we have an equilibrium story to tell. In particular, on the equilibrium path we will find that $p = P$, so that $p = \sqrt{2} - 1$.

¹²I hope that it is uncontroversial to suggest that Selten's theory of the "trembling hand" is too simple a story to be appropriate in the current context.

¹³For example, on his n th entry into the information set I he might perhaps arbitrarily estimate the probability that he has yet to reach node d_1 as the discounted sum

$$Q_n = (1 - \Delta)^{-1} (q_{n-1} + \Delta q_{n-2} + \Delta^2 q_{n-3} + \dots),$$

where, for each $k < n$, $q_k = p_k/(1 + p_k)$ and p_k is the probability with which he actually went straight ahead at intersections during his k th entry into I . On the assumption that Q_n is correct, he can then choose p_n optimally to be $(1 - 2Q_n)/(1 + 2Q_n)$. Then q_n is defined in terms of its predecessor by $q_n = p_n/(1 + p_n)$. The question is then whether the sequence $\{q_n\}$ converges.

5 A Team Analysis

My doubts about modeling players as teams of agents with identical preferences have already been mentioned. But it cannot be denied that the methodology has its advantages in games of perfect recall in which no player ever moves more than once. In the presence of the latter proviso, it is irrelevant what would be inferred about a player's future play from a counterfactual deviation on his part from the equilibrium path. At the same time, the "single improvement property" for games of perfect recall ensures that there is no loss of efficiency in allowing agents to make decisions one-by-one (Piccione and Rubinstein [12]).

Gilboa [6] offers the orthodox case for modeling the the absent-minded driver as a team problem. The driver is treated as though he had a multiple personality, with two personalities or agents to be called Alice and Bob. They have the same preferences but different information. When one of the two agents is called upon to make a decision, the agent does not know whether he is at the first exit or the second, but he does know that the other agent will be making an independent decision at the other exit (should it be reached). Figure 2(a) shows a representation of the game of perfect recall that must then be played between Alice and Bob. Its opening move is a chance move that assigns control of the driver at the first exit to Alice or Bob with equal probability.¹⁴

Figure 2: Splitting the driver's personality

Figure 2(b) shows the strategic form of the game. Its unique equilibrium

¹⁴Although nothing in the specification of the Paradox of the Absent-Minded Driver would seem to provide a strong justification for setting the probability, τ with which the opening chance move assigns control to Alice equal to $1/2$. But if $\tau \neq 1/2$, we shall not be led to the satisfying conclusion that Alice and Bob should each choose R or S with probability $1/2$.

is easily calculated. Alice and Bob should each independently choose R or S with probability $\frac{1}{2}$. The outcome is then the same as when the driver committed himself to the behavioral strategy $b(\frac{1}{2})$ before reaching I in the analysis of Piccione and Rubinstein [12].¹⁵ In particular, both Alice and Bob receive an expected payoff of $\frac{1}{4}$.

In spite of the welcome conclusion to which this analysis leads, I believe it to be wrong, because it neglects the fact that what a person does in certain circumstances must surely be evidence about what he will do if placed in exactly same circumstances later on. It is certainly true that arguments that take account of such reasoning can easily go astray. Somehow a line must be drawn between valid uses of the principle and invalid uses that lead authors like Nozick [11] to twinning arguments which supposedly demonstrate that cooperation is rational in the Prisoners Dilemma (Binmore [5, p.205]). As in much else, my own view is that the right way or ways to proceed will remain mysterious until we have satisfactory *algorithmic* models of the players we study (Binmore [4]).

6 Conclusion

This paper has done no more than comment on the comments made by Piccione and Rubinstein [12] on their Paradox of the Absent-Minded Driver. It claims that the paradox is illusory when the driver has frequency data to support his beliefs. But, as Gilboa [6] aptly observes, what counts as a paradox depends on the viewpoint of the observer. As for the paradox in the general case, I have nothing useful to say beyond the observation that it highlights the need for an adequate theory of decision-making under uncertainty to supplement the current Bayesian orthodoxy.

References

- [1] S. Alpern. Games with repeated decisions. *SIAM Journal of Control and Optimization*, 468-477, 1988.
- [2] K. Binmore. Debayesing game theory. In B. Skyrms, editor, *Studies in Logic and the Foundations of Game Theory: Proceedings of the Ninth International Congress of Logic, Methodology and the Philosophy of Science*, Kluwer, Dordrecht, 1992.
- [3] K. Binmore. *Fun and Games*. D. C. Heath, Lexington, Mass., 1991.
- [4] K. Binmore. Modeling rational players, I and H. *Economics and Philosophy*, 3 and 4:179-214 and 9-55, 1987.

¹⁵This is not an accident. Whatever payoffs are assigned to the possible outcomes in the game, the two analysis yield the same result.

- [5] K. Binmore. *Playing Fair: Game Theory and the Social Contract I*. MIT Press, Cambridge, Mass., 1994.
- [6] I. Gilboa. Technical Report.
- [7] J. Isbell. Finitary games. In *Contributions to the Theory of Games III*, Princeton University Press, Princeton, 1957.
- [8] J. M. Keynes. *A Treatise on Probability*. Macmillan, London, 1921.
- [9] H. Kuhn. Extensive games and the problem of information. In H. Kuhn and A. Tucker, editors, *Contributions to the Theory of Games II*, Princeton University Press, Princeton, 1953.
- [10] E. McClennen. *Rationality and Dynamic Choice*. Cambridge University Press, Cambridge, 1990,
- [11] R. Nozick. Newcomb's problem and two principles of choice. In N. Rescher, editor, *Essays in Honor of Carl G. Hempel*, Reidel, Dordrecht, Netherlands, 1969.
- [12] M. Piccione and A. Rubinstein. On the interpretation of decision problems with imperfect recall. 1994. Working Paper 9-94, Sackler Institute, Tel Aviv University.
- [13] A. Rubinstein. Finite automata play the repeated prisoners' dilemma. *Journal of Economic Theory*, 39:83-96, 1986.
- [14] L. Savage. *The Foundations of Statistics*. Wiley, New York, 1951,
- [15] R. Selten. Reexamination of the perfectness concept for equilibrium points in extensive-games. *International Journal of Game Theory*, 4:25-55, 1975.
- [16] R. Strotz. Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23:166-180, 1973.
- [17] J. Von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1944.