GAMES AND ECONOMIC BEHAVIOR **20,** 121–130 (1997)
ARTICLE NO. GA970579

The Absent-Minded Driver's Paradox: Synthesis and Responses*

Michele Piccione

Department of Economics, University of British Columbia, Canada

and

Ariel Rubinstein

School of Economics, Tel Aviv University, Israel and
Department of Economics, Princeton University, Princeton, New Jersey

Received April 18, 1997

We classify the responses to our paper "On the Interpretation of Decision Problems with Imperfect Recall" and address some of the points raised. © 1997 Academic Press

1. THE AIM OF OUR ORIGINAL PAPER

The aim of our original paper, Piccione and Rubinstein (1997) (P & R from now on), was to point out that the model commonly used to describe a decision problem with imperfect recall suffers from major ambiguities in its interpretation. We claimed that several issues which were immaterial in decision problems with perfect recall may be of importance in the analysis of decision problems with imperfect recall. The issues that we raised can be summarized by the following questions:

1. What decisions can be made? In particular, can a decision maker decide about when to make a decision?

2. What is the timing of decisions? Is there a planning stage or are decisions made only at the time actions are executed?

* We are grateful to Northwestern University and New York University for their hospitality during the period in which this research has been conducted. We thank many friends, colleagues, readers, and listeners with whom we discussed the issues raised in our paper. The list is very long. We respect their criticism and are grateful for the encouragement.

0899-8256/97 $25.00

3. Can a decision maker change his strategy along its execution? And, if he does change his strategy, can he change it again?

4. Can a decision maker use random devices?

The answers for these questions are not specified by the conventional interpretation of the model of extensive decision problems. Indeed, these issues are immaterial for the analysis of optimal (rational) behavior in problems with perfect recall. The objective of P & R was: (i) to demonstrate that these issues are crucial for the analysis of some decision problems with imperfect recall and (ii) to draw the boundaries around families of problems for which the issues are indeed immaterial.

The analysis of decision problems with the special form of imperfect recall which we named ‘‘absent-mindedness’’ was particularly problematic. Several issues were uncovered in the example we called the ‘‘absent-minded driver’s paradox.’’ In our opinion, the driver’s example has a paradoxical flavor due to the conflict between two ways of reasoning at an intersection. The first instructs the decision maker to follow his initial decision not to exit, following an intuitive principle of rationality that unless new information is received or there is a change in tastes, previous decisions should not be changed. The second way of reasoning, maximizing the expected payoff given the belief, suggests he should deviate from his initial decision.

We wish to emphasize that this example was *not* the center of our paper. Nevertheless, as we expected, it received most of the attention. As Lipman points out, this example ‘‘touched a nerve.’’ There is always something annoying in any paradoxical example, perhaps because of the deceptive way by which paradoxes are framed. It is our impression that we touched two delicate and related issues: first, the wide-spread sentiment among game theorists that the model of extensive games is the ‘‘correct’’ model to express any sort of strategic situations; second, the wide-spread practice among economists of dispensing with procedural aspects of decision making.

Of course, we are honored that, in Lipman’s words, ‘‘so many eminent and intelligent people’’ have responded to our paper, provided fair criticism and new insights into the topic of decision making with imperfect recall. Naturally, we wish to reply to their comments. First, however, we think it best to help readers to find their way among the various papers.

2. SYNTHESIS

The papers in this symposium can be classified, in our opinion, according to two interpretative issues raised in P & R:

(1) Is there a preplay planning stage?

(2) At a point in time, can a decision maker control his behavior for later instances he might encounter?

In the mixed strategy version of the paradox in P & R, the conclusion reached in the preplay planning stage (the optimal strategy) is reviewed by the driver at the intersection from the perspective that he can control his future behavior.

Battigalli (1997) assumes that there is preplay planning but a decision maker is allowed to control his future behavior only for information sets not reachable under his planned strategy.

Halpern (1997) (in his notion of gt-consistency) assumes that there is preplay planning and that replanning at an information set occurs only the first time the information set is reached (the first intersection in the driver's example).

Aumann, Hart, and Perry (1997a) assume that there is a planning stage, but argue in favor of what we called the modified multiselves approach; a player can only control his current action and considers the rest of his play as fixed.

Gilboa (1997) does not allow for a planning stage and assumes that a decision maker has no control beyond the instance in which he operates. The driver is viewed as a collection of two agents who take one another's behavior as given.

For the absent-minded driver's problem, all these papers provide alternative frameworks which yield the "exit with probability 1/3" strategy as the consistent rule of behavior. Lipman describes an additional approach in which at an intersection the decision maker assumes that the (full) strategy he picks will be followed at the immediate instance and at any other future instance. In the driver's example, the only strategy which is consistent with this assumption is "exiting with probability 5/9."

Other issues were also addressed. Grove and Halpern tackle the question of belief formation to compare ex-ante and ex-post payoffs of a fixed strategy. They observe that, in the case of absentmindedness, reaching a node in the information set changes the initial assessment of the probability of reaching the terminal nodes. They refer to this as the "probabilistic version of the paradox." *Aumann, Hart, and Perry* (1997b) reiterate this point and show that this phenomenon occurs in general imperfect recall situations.

An important new element appears in the last part of *Halpern*'s paper. Because information sets do not capture possible restrictions on the knowledge that a decision maker has of the strategy in use, Halpern suggests a new model in which statements about this type of knowledge can be made explicitly.

We now turn to a more detailed examination of these papers. For views on these issues not included in this symposium, we refer the reader to Dekel and Gul (1996), Matsui (1996), and Segal (1995).

3. RESPONSES

(a) *Battigalli*

Between the assumption that a decision maker fully controls his future behavior and the other extreme one that he controls only the current action there is a wide scope of interim options. Battigalli's paper suggests one such possibility. The main assumption is that a player who reconsiders his plan at an information set controls (i) his action at that information set (only at that instance) and (ii) his behavior at the future information sets which cannot be reached with positive probability according to the original plan. The rationale behind this assumption is that a player cannot instruct a deviation from the initial plan to the agents who control an information set reachable under the initial plan. Referring to Example 2 in P & R, Battigalli says: "How can the decision maker at information set d_1 plan to change her behavior at d_3, where a deviation from the original plan cannot be observed?" According to Battigalli an agent can receive "orders" to deviate from a plan only if the plan did not call for him to be active. This leads the definition of "constrained time consistent" strategies. These are strategies which, at any information set along their execution, are immune to deviations in the current action (without affecting the action to be taken at that information set should it be revisited) and in the actions at information sets which are not reachable under the original strategy.

Battigalli's main results are (1) any optimal strategy is constrained time consistent and (2) every constrained time consistent strategy (and not only optimal strategy) is modified multiself-consistent. Constrained time consistent strategies are not necessarily optimal; this is similar to time and modified multiself-consistent strategies.

Battigalli's concept demonstrates the richness of the family of plausible consistency assumptions which can be made. The paper reinforces our view that one needs to enrich the model to determine the appropriate consistency requirements. We agree with Battigalli that his concept is appropriate for cases where the player is a team and communication difficulties result in lack of perfect information. Constraints in transferring information, however, do not necessarily imply inability to transfer "orders" as to the action to be taken. Consider for example a battlefield situation (which could be the story behind Example 2 in P & R) in which an officer observes the arena but the officer in charge of artillery does not. It may be that the officer in the field is unable to describe the enemy's position to the artillery officer, whereas the officer can send him a message as to whether to open fire. Analogously, it might be that a decision maker receives at the outset a very detailed situation report which he is unable to memorize whereas he can keep in mind the conclusion about what to do at the second date.

(b) *Halpern*

Halpern's paper consists of two parts. In the first part, the author reviews four types of time consistency. One of his innovative contributions is the definition of gt-time consistency: a decision maker reviewing his strategy at an information set believes that this review is done at the first node. The author identifies a class of decision problems (of "partial recall," which includes the absent-minded driver example) for which, under the assumption that beliefs are concentrated on the upper contour of an information set, optimality is equivalent to time consistency.

We find it hard to make sense of this assumption. One possible interpretation is that the decision maker infers from his desire to change the plan that it must be the first time the information set has been reached since changes would be made the first time the opportunity arises. It is not clear, however, why a decision maker who draws inferences from his desire to change the strategy should not also make inferences from his unwillingness to change it. When the original plan is confirmed, why should his inference exclude the possibility of being at the second intersection?

In the second part of the paper, the author proposes an alternative model for decision problems with imperfect recall. This model contains some of the elements Fagin *et al*. (1995) introduced in the context of computer sciences. The central contribution is the suggestion of a language which allows statements about the decision maker's "knowledge" of the strategy used.

(c) *Gilboa*

Gilboa excludes any form of preplanning. In his opinion, the right formulation of the driver's problem is a game in which the players are two independent agents of the decision maker. He states "decision problems can and should be formulated in such a way that information sets cannot contain more than one decision node on each path." The decision maker's problem is identified with a two-player game in which no agent controls the other. The game has a unique symmetric equilibrium which coincides with the driver's optimal strategy.

We are not convinced by Gilboa's unreserved statement that "the language of agents is appropriate whenever a dynamic game is concerned." By definition, this language eliminates the circumstance of the driver controlling his actions at a future intersection! The basic formalization of the problem defuses any paradoxical aspects.

Our major criticism of Gilboa's approach rests on his claim of ultimate "correctness" of the two-agent game. Most game theorists think of a game as a physical description of a situation. Alternatively, a game can be thought of as a presentation of the way in which a situation is perceived by the players (see, for example, Rubinstein, 1991). Gilboa goes much further and, in his own words, analyzes the situation in "terms which can only be "observed" by introspection." In our opinion, the game he suggests fits neither any physical description nor any reasonable perception of the situation. We find it hard to believe that anybody would think about the driver's example from the point of view of having two selves equally likely to be called to act first and each taking the other as absolutely given.

In principle, Gilboa's game allows for the possibility that an agent believes that his "twin agent" may behave according to a different mode of behavior. Indeed, the game has also an asymmetric equilibrium in which one agent exits and the other continues. This equilibrium has a payoff of 2 which cannot be achieved in the original decision problem.

(d) *Aumann-Hart-Perry*

The main message of *Aumann-Hart-Perry* (1997a) is a strong endorsement of the view that the *only* correct approach to the problem is the modified multiselves approach (which they call "action-optimality"). Their main claim is that: "when at one intersection, he (the driver) can determine the *action only there*, and not at the other intersection—where he isn't" and "whatever reasoning obtains at one must obtain also at the other; and, he is aware of this." In their belief, controlling the action at the future intersection can only be done by wizardry. In response to the fact that the multiagents approach may not lead to optimality, the authors add that "because of his absent-mindedness, . . . , coordination can take place only before he starts out-at the planning stage." These are possible and even plausible assumptions but we fail to see any compelling reasons for making these assumptions the *only* understanding of an extensive game. For further discussion of this point and the connection between the Gilboa and Aumann, Hart, and Perry papers see Lipman.

As to the paradox, the authors say that "one can imagine scenarios for which these observations do not hold" and add "P & R have yet to adduce an explicit scenario that does display a paradox." We do not see what is meant by an "explicit scenario" in this context. Aumann, Hart, and Perry would probably agree that it would be quite a feat to produce an "explicit" barber who shaves all those and only those who do not shave themselves. If the scenario must include details about the way that the decision maker reasons, we have one. If they require that the process of reasoning be derived from primitive assumptions which induce consistent behavior, well, obviously such a scenario, by definition, will not display a paradox.

(e) *Grove and Halpern*

Grove and Halpern address the issue of beliefs at an information set in games with absent-mindedness. In particular, they attempt to resolve what they call the "expectation paradox." In the example of the absent-minded driver, the ex-ante expected value of strategy is $p^2 + 4p(1-p)$, where p is the probability of continuing. Conditional upon being at the information set, the expected payoff is $\alpha[p^2 + 4p(1-p)] + (1-\alpha)[p + 4(1-p)]$, where α is the belief that the decision maker assigns to being at the first node. The ex-ante payoff and the payoff at the information set differ unless α is equal to 1. For the case of the ex-ante optimal strategy, Grove and Halpern observe that "the driver's ex-ante valuation of the game is $4/3$, but as soon as he reaches an exit (which he knew was certain to happen) he thinks the game is worth more. He would refuse an offer of $4/3$ to quit the game once it had started, yet at the bar would believe this to be a fair offer."

To tackle this problem, the authors define a state space which includes all (z, x) such that z is a final history and x is a node on the path to z, and a probability measure $q_b(z, x)$ on this space. $q_b(z, x)$ is the probability of being on the path to z and visiting the history x under strategy b.

The authors compare two probabilistic models. Given a behavioral strategy b, let $p_b(z)$ be the probability of reaching z defined in the standard way:

(1) $q_b(z, x)$ is proportional to $p_b(z)/$[number of times z crosses the information set that x is in]. This measure corresponds to our definition of z-consistency. The authors observe that under this measure, if the utility at (z, x) is defined to be the utility at z, the ex-ante expected utility of a strategy is identical to the expected utility conditional upon being at the information set.

(2) $q_b(z, x) = p_b(z)/\Sigma_{(z', x')} p_b(z')$. In this case, the induced beliefs are the frequency probabilities. However, ex-ante expected utility differs from the expected utility at the information set.

The comparison of ex-ante and ex-post payoffs of a fixed strategy is the sole objective of Grove and Halpern. For a discussion of this issue we refer the readers to Lipman.

4. THE PARADOX

Our Position

The word "paradox" is often a catalyst for heated discussions. We would be happy to make do with an "interesting" or "provocative and worth having" example as Aumann, Hart, and Perry kindly called it. Despite all that has been said we are still confused about its resolution.

We do not find it surprising that the paradox disappears when interpretative ambiguities are removed by making specific assumptions. However, the suggestions by different eminent scholars are far from coinciding. Some authors also provide arguments (see Binmore, 1996 and Lipman, 1997) in support of the reasonableness of a strategy different from the ex-ante optimal one.

Contrary to what Aumann, Hart, and Perry speculated, we *do not agree* "that the pure strategy case is not particularly interesting" and do think that the paradoxical flavor of the example appears in full strength in this case. Rationality is always defined with respect to choice, *given* a set of alternatives. We fail to see why, to define rational behavior, the set of alternatives *must* (of course, it may) include the possibility of using randomizing devices. The availability of random elements is not at all obvious.

Another Variant of the Paradox

Suppose that we have a procedure which computes the "rational" behavioral strategy for all decision problems with the form of the driver's example. The input is the payoffs of the three possible consequences and the output is a lottery between "exit" and "continue." This program is activated by the driver to determine his strategy at each intersection.

Several tests of "rationality" can be constructed for this program. One is the following. After the output is produced, a line is added: Type "Y" to execute the outcome of the program. Type "N" to erase the program and to replace it with a program which automatically executes the strategy "continue with probability p^*" when an intersection is reached.

At an intuitive level, "rationality" would imply that the decision maker should type "Y" to the question independently of p^*. However, if the output of the program is the ex-ante optimal strategy and the decision maker believes he would have acted "rationally" typing Y at previous occurrences, then typing Y now conflicts with expected utility maximization for some values of p^*.

The only strategy for which the answer Y is always consistent with expected utility maximization is the "time consistent strategy" to continue probability $5/9$. Of course, it is easy to find a "test" which this strategy fails to pass. This reinforces our view as to the presence of interpretative ambiguities in the model.

5. THE SEARCH FOR ONE TRUTH

The main goal of our paper was to illustrate that modelling aspects which are immaterial in decision problems with perfect recall are of great importance in decision problems with imperfect recall. We believe we have a consensus on this issue. Details such as strategy recall, circumstances which induce a decision maker to reassess his strategy, and so on, are recognized to be crucial for the analysis of such decision problems. There is less of a consensus about the next step. Some believe that the standard model of extensive games can tackle these issues successfully. Others think that there is a need to expand the model to allow explicit consideration of the above elements. The only paper in this symposium which attempts to enrich the model in this direction is Halpern's.

Lipman objects to the view "that there is one true way to model the perceptions of an agent with imperfect recall." We fully agree with his criticism. We see no justification for ruling out *a priori* that a decision maker could consider a change in his plan beyond the instance of the deliberation. Calling it a "mysterious psychic process" or an "incorrect assumption" (see Aumann, Hart, and Perry, 1997a, 1997b) is not making it more compelling. Every point of view can be criticized in such terms. Take the "modified multiselves approach"; what "mysterious psychic process" leads a decision maker to believe that a change in the current action is only a one shot event and is no evidence about future behavior, despite the fact that identical circumstances will appear later on? Or, more generally, what "mysterious psychic process" validates that a rational decision maker behaves as if composed of independent agents?

We agree with Binmore in that "the right way or ways to proceed will remain mysterious until we have satisfactory algorithmic models of the players we study." In our opinion, decision theory and game theory cannot give sharp answers. Decision problems such as the driver's example can be reasonably analyzed in a variety of ways. As formal theoreticians we can at most clarify some of the logic of these ways.

REFERENCES

Aumann, R., Hart, S., and Perry, M. (1997a). "The Absent-Minded Driver," *Games Econ. Behav.* **20**, 102–116.

Aumann, R., Hart, S., and Perry, M. (1997b). "The Forgetful Passenger," *Games Econ. Behav.* **20**, 117–120.

Battigalli, P. (1997). "Dynamic Consistency and Imperfect Recall," *Games Econ. Behav.* **20**, 31–50.

Binmore, K. (1996). "A Note on Imperfect Recall," in *Understanding Strategic Interaction—Essays in Honor of Reinhard Selten* (W. Guth *et al.*, Ed.), pp. 51–62, Springer-Verlag, New York/Berlin.

Dekel, E., and Gul, F. (1996). "Rationality and Knowledge in Game Theory," in *Advances in Economic Theory, Seventh World Congress* (D. Kreps and K. F. Wallis, Eds.), pp. 87–173, Cambridge Univ. Press, Cambridge.

Fagin, R., Halpern, J., Moses, Y., and Vardi, M. (1995). *Reasoning about Knowledge,* MIT Press, Cambridge, MA.

Gilboa, I. (1997). "A Comment on the Absent Minded Driver Paradox," *Games Econ. Behav.* **20**, 25–30.

Grove, A., and Halpern, J. (1997). "On the Expected Value of Games with Absentmindedness," *Games Econ. Behav.* **20**, 51–65.

Halpern, J. (1997). "On Ambiguities in the Interpretation of Games Trees," *Games Econ. Behav.* **20**.

Lipman, B. (1997). "More Absentmindedness," *Games Econ. Behav.* **20**, 97–101.

Matsui, A. (1996). "Multiple Minds in Game Theory," mimeo.

Piccione, M., and Rubinstein, A. (1997). "On the Interpretation of Decision Problems with Imperfect Recall," *Games Econ. Behav.* **20**, 3–24.

Rubinstein, A. (1991). "Comments on the Interpretation of Game Theory," *Econometrica* **59**, 909–924.

Segal, U. (1995). "Don't Fool Yourself to Believe You Won't Fool Yourself Again," mimeo.