

The Absent-Minded Driver*

Robert J. Aumann, Sergiu Hart, and Motty Perry

*Center for Rationality and Interactive Decision Theory, The Hebrew University of
Jerusalem, Feldman Building, Givat-Ram, 91904 Jerusalem, Israel*[†]

Received February 6, 1996

The example of the “absent-minded driver” was introduced by Piccione and Rubinstein in the context of games and decision problems with imperfect recall. They claim that a “paradox” or “inconsistency” arises when the decision reached at the “planning stage” is compared with that at the “action stage.” Though the example is provocative and worth having, their analysis is questionable. A careful analysis reveals that while the considerations at the planning and action stages do differ, there is no paradox or inconsistency. *Journal of Economic Literature Classification Numbers: D81, C72.* © 1997 Academic Press

1. INTRODUCTION

An absent-minded driver starts driving at START in Figure 1. At X he can either EXIT and get to A (for a payoff of 0) or CONTINUE to Y . At Y he can either EXIT and get to B (payoff 4), or CONTINUE to C (payoff 1). The essential assumption is that he cannot distinguish between intersections X and Y , and cannot remember whether he has already gone through one of them.

Piccione and Rubinstein (1997; henceforth P & R), who introduced this example, claim that a “paradox” or “inconsistency” arises when the decision reached at the *planning stage*—at START—is compared with that at the *action stage*—when the driver is at an intersection. Though the example is provocative and worth having, P & R’s analysis seems flawed. A careful analysis reveals that while the considerations at the planning and action stages do differ, there is no paradox or inconsistency.

*This is an outgrowth of notes and correspondence originating in May and June of 1994. We thank Ehud Kalai, Roger Myerson, Phil Reny, Ariel Rubinstein, Dov Samet, Larry Samuelson, and Asher Wolinsky for discussions on these topics, and the Associate Editor and the referee for some insightful comments. Research partially supported by grants of the U.S.–Israel Binational Science Foundation, the Israel Academy of Sciences and Humanities, and the Game Theory Program at SUNY–Stony Brook.

[†]E-mail: ratio@vms.huji.ac.il.




FIG. 1. The absent-minded driver problem.

We start in Section 2 by laying down the fundamental observations that underlie the driver's decision problem, and then show in Section 3 how P & R's analysis violates these observations. In Section 4 we formally define the concept of *action-optimality* and use it to analyze P & R's example. Section 5 studies action-optimality in a more general setup, with some interesting and unexpected conclusions. We conclude with a detailed discussion of various issues in Section 6.

2. FUNDAMENTALS

At the **planning** stage, the decision problem is straightforward. In the example, the optimal (randomized) decision is¹ "CONTINUE with probability 2/3 and EXIT with probability 1/3." We call this the *planning-optimal* decision.

¹The problem is to maximize $(1 - p) \cdot 0 + p(1 - p) \cdot 4 + p^2 \cdot 1$ over p , where p is the probability of CONTINUE.

At the **action** stage, though, even *formulating* the decision problem is not straightforward. The following observations are essential for a correct analysis of the decision at the action stage.²

- **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 isn't.

- **Second**, since he is in completely indistinguishable situations at the two intersections, whatever reasoning obtains at one must obtain also at the other, and he is aware of this.

3. THE P & R ANALYSIS

Consider the action stage. The driver finds himself at an intersection; he does not know which. Let α be the probability that X is the current intersection, and let p and q be the probabilities of CONTINUE at the current and at “other” intersections, respectively. Then the expected payoff at the action stage is

$$H(p, q, \alpha) := \alpha[(1 - p) \cdot 0 + p(1 - q) \cdot 4 + pq \cdot 1] \\ + (1 - \alpha)[(1 - p) \cdot 4 + p \cdot 1].$$

P & R maximize $H(p, p, \alpha) = \alpha[(1 - p) \cdot 0 + p(1 - p) \cdot 4 + p^2 \cdot 1] + (1 - \alpha)[(1 - p) \cdot 4 + p \cdot 1]$ over p , holding α fixed. Thus they take p and q as decision variables to be maximized simultaneously, subject to the constraint $q = p$. This makes sense only if the driver controls the probabilities at both intersections—a violation of the first observation. But even if, by some magical process, the driver *could* control the probability q at the other intersection, surely α depends on q , and cannot be held fixed in the maximization!

4. ACTION-OPTIMALITY

How, then, *should* the driver reason at the action stage? Let us spell out in detail the implications of the two observations in Section 2:

- (i) The optimal decision is the same at both intersections; call it p^* .

²One can imagine scenarios for which these observations do not hold. But such scenarios do not correspond to the plain meaning of the words used to describe the situation. More important, with those other scenarios the analysis at the planning stage also changes, and again there is no paradox. Piccione and Rubinstein have yet to adduce an explicit scenario that *does* display a paradox. See Section 6(c) for more on this issue.

(ii) Therefore, at each intersection, the driver believes that p^* is chosen at the other intersection.

(iii) At each intersection, the driver optimizes his decision given his beliefs. Therefore, choosing p at the current intersection to be p^* must be optimal given the belief that p^* is chosen at the other intersection.

Given a behavior q at the other intersection, the probability that the current intersection is X is³ $\alpha = 1/(1 + q)$. Thus if we set $h(p, q) := H(p, q, 1/(1 + q))$, we can restate the final implication (iii) as follows:

p^* is action-optimal if the maximum of $h(p, p^*)$ over p
is attained at $p = p^*$.

Thus, p^* is a fixed point⁴ of the (set-valued) mapping $q \rightarrow \arg \max_p h(p, q)$.

Applying this analysis to the example, we see that the (randomized⁵) planning-optimal decision—CONTINUE with probability $2/3$ —is also action-optimal. Indeed, if this is the behavior at the other intersection, then the probability that the current intersection is X is $\alpha = 3/5$. Therefore the expected payoff from choosing CONTINUE at the current intersection with probability p is $h(p, 2/3) = (3/5) \cdot [(1 - p) \cdot 0 + p \cdot (1/3) \cdot 4 + p \cdot (2/3) \cdot 1] + (2/5) \cdot [(1 - p) \cdot 4 + p \cdot 1]$, which equals $8/5$ (for all p). So $p = 2/3$ maximizes it; thus $p^* = 2/3$ is action-optimal.

So there is no paradox: the planning-optimal choice of $2/3$ is also action-optimal.

Moreover, $p^* = 2/3$ is the *unique* action-optimal decision. Indeed,

$$\begin{aligned} h(p, q) &= \frac{1}{1 + q} [(1 - p) \cdot 0 + p(1 - q) \cdot 4 + pq \cdot 1] \\ &\quad + \frac{q}{1 + q} [(1 - p) \cdot 4 + p \cdot 1] \\ &= \frac{(4 - 6q)p + 4q}{1 + q}. \end{aligned}$$

³This probability is the “consistent belief” of P & R; but unlike P & R, we consider it to be *the* one that is appropriate to this problem. In the Appendix we derive it from a formal model. Informally, since the driver always goes through X , but only q of the time through Y , the ratio of the probabilities is 1 to q , so they must be $1/(1 + q)$ and $q/(1 + q)$. These probabilities may also be derived from a “fair lottery” approach (see Footnote 3 in Aumann, Hart, and Perry, 1997).

⁴Formally, (p^*, p^*) 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 game with payoff functions h).

⁵For the pure case, see Section 6(d) below.




FIG. 2. Multiple action-optimal decisions.

Given q , the maximizing p therefore is: 0 for $q > 2/3$; 1 for $q < 2/3$; and anything for $q = 2/3$. Thus the only fixed point is $p^* = 2/3$.

The notion of action-optimality defined here is mathematically identical to the “modified multiselves approach” described near the end (Section 7) of P & R. But unlike P & R, we consider this notion to be *the* natural and correct formulation of the driver’s decision problem at the action stage. See Section 6 below for further discussion.

5. A MORE CHALLENGING EXAMPLE, WITH MULTIPLE ACTION-OPTIMAL DECISIONS

We have seen that, in the specific example of P & R, randomized planning optimality and action optimality coincide. This is not always so! While, in general, any planning-optimal randomized decision is also action-optimal,⁶ we will now show that there may be action-optimal choices that are not planning-optimal. Surprisingly, there may even be action-optimal choices that, at the intersections, look better than the planning-optimal choice.

⁶Proposition 3 of P & R and also the end of the Appendix below.




FIG. 3. A more challenging example.

For a simple example with action-optimal decisions that are not planning-optimal, change the payoffs to be 1 at A , 0 at B , and 2 at C —see Fig. 2. The unique planning-optimal choice is CONTINUE, i.e., $\bar{p} = 1$. There are, however, three action-optimal decisions: $p_1^* = \bar{p} = 1$, $p_2^* = 0$, and $p_3^* = 1/4$ (e.g., to see that $p_2^* = 0$ is action-optimal, note that if the decision at the other intersection is EXIT, then it is indeed optimal to EXIT now too).

This leads us to the next point: When there are multiple action-optimal choices, how do their payoffs compare? Of course, when computed at START, the one that is planning-optimal yields the maximum payoff. But how does it look at the current intersection? In the example above, $p_1^* = \bar{p} = 1$ yields the highest payoff among all the action-optimal decisions also when compared at the intersections; i.e., $h(p_1^*, p_1^*) > h(p_i^*, p_i^*)$ for $i = 2, 3$. But this is not always so when there are more than two intersections. Indeed, consider an example with three intersections, the

payoff being 7 at the first EXIT, 0 at the second EXIT, 22 at the third EXIT, and 2 if always CONTINUE (see Fig. 3). Then:

(i) The unique planning-optimal decision is $\bar{p} = 0$ (with a payoff of 7).

(ii) There are three action-optimal decisions: $p_1^* = \bar{p} = 0$, $p_2^* = 7/30$, and $p_3^* = 1/2$.

(iii) The ex-ante expected payoffs for p_1^* , p_2^* , and p_3^* are, respectively, 7, $8519/1350 \approx 6.31$, and 6.5.

(iv) The ex-post expected payoffs $h(p^*, p^*)$ for p_1^* , p_2^* , and p_3^* are, respectively, 7, $7378/1159 \approx 6.37$, and $50/7 \approx 7.14$; thus $h(p_3^*, p_3^*)$ is larger than $h(p_1^*, p_1^*) = h(\bar{p}, \bar{p})$.

The reader may ask, since the choice is among three possibilities yielding 7, ≈ 6.37 , or ≈ 7.14 , why does the driver not choose the action with the highest yield, namely p_3^* ? The answer, of course, is that at the action stage, the driver *cannot* choose among p_1^* , p_2^* , and p_3^* . His beliefs there are not under his control; he cannot choose what to believe. Action optimality is a condition for consistency of beliefs and rational behavior. If the player is to be consistent at the action stage, he must believe in one of the three possibilities p_1^* , p_2^* , p_3^* ; but *which* one is not up to him at that stage.

We have already pointed to the formal similarity between action-optimality and game equilibria. Choosing *between* the p_i^* is much like choosing between game equilibria, which is something that the individual player in a game cannot do—it must be done by an outside force (like a custom or a norm), or by all the players somehow coordinating their actions.

In our case, there is only one player, who acts at different times. Because of his absent-mindedness, he had better coordinate his actions; this coordination can take place only before he starts out—at the planning stage. At that point, he should choose p_1^* . If indeed he chose p_1^* , there is no problem. If by mistake he chose p_2^* or p_3^* , then that is what he should do at the action stage. (If he chose something else, or nothing at all, then at the action stage he will have some hard thinking to do.)

Once having coordinated on p_1^* , there is no incentive for the driver to choose p_3^* at the action stage. Nevertheless, one may ask whether at that stage he will be sorry that he did not coordinate on p_3^* rather than on p_1^* . After all, if he had, his expectation now would be ≈ 7.14 , which is greater than the 7 he now expects! If the answer is “yes,” we have a conceptual puzzle; why should the driver get himself into a situation where at START he is *sure* that he will be sorry at every intersection he reaches? Why not avoid the sorrow by coordinating on p_3^* in the first place?




FIG. 4. An automatic car.

But the answer is “no”; he should not be sorry. Having chosen p_1^* , he knows he must be at X when finding himself at an intersection. Being at X is like being at $START$ —i.e., at the planning stage—and then the best choice is p_1^* and *only* p_1^* . So when reaching an intersection after having chosen p_1^* , the driver is *not* sorry that he indeed chose p_1^* rather than p_3^* .

Why, then, is the driver’s expectation at an intersection nevertheless larger for p_3^* than for p_1^* ? The reason is that at an intersection, his belief as to where he is if he chose p_3^* differs from his belief when he chose p_1^* . Having chosen p_1^* , he knows he must be at X . If he had chosen p_3^* , he would have attributed probabilities $4/7$, $2/7$, and $1/7$ to being at X , Y , and Z . He “prefers” the latter distribution, because it gives him a chance of already having passed the “dangerous” intersection Y and a better shot at the high payoff of 22. But as we said above, one cannot choose one’s beliefs, and it makes little sense to discuss “preferences” between them. Specifically, since he does know that he is at X , it would be silly for him to say, “I wish I had chosen the other plan, because then in my ignorance I would have been deluded into expecting a higher payoff than now.”

To clarify this point, consider the example in Fig. 4. The car is automatic and $EXITS$ with probability $1/2$ at each intersection. The decision maker is




FIG. 5. Clear-headed or absent-minded?

a passenger, who sleeps during most of the trip. At *START*, he is given the option to be woken either at both intersections, or only at *X*. In the first option he is absent-minded: when waking up, he does not know at which intersection he is. We call the second option “clear-headedness.”

As in the previous discussion, the question at *X* is not operative—what to do—but only whether it makes sense to be “sorry.” If he chose clear-headedness, his expectation upon reaching *X* is $1/4$. If he had chosen absent-mindedness, then when reaching *X* he would have attributed probability $2/3$ to being at *X* and $1/3$ to being at *Y*. Therefore his expected payoff at that point would have been $(2/3) \cdot (1/4) + (1/3) \cdot (1/2) = 1/3$, which is larger than $1/4$. Is he therefore sorry that he chose to be clear-headed? Clearly this would be absurd, as the payoffs do not depend on his choice.

To make this even more striking, assume that when he is not absent-minded, the probability of *CONTINUE* is increased to $4/7$ (Fig. 5). Then ex-ante the clear-headed option is actually preferred to the absent-minded option—it yields $16/49$ rather than $1/4$. But, upon reaching *X*, the clear-headed option still yields $16/49$, whereas the absent-minded option

yields $1/3$ (as above), which is bigger than $16/49$. Surely, it would be absurd for the decision maker to wish he were absent-minded—it would be sticking his head in the sand!

Another issue that these examples raise is this⁷: Assume the driver in Fig. 3 has chosen p_3^* at START. His expected payoff there is then 6.5. But, he is *certain* to go through X , where his expected payoff becomes ≈ 7.14 . For a treatment of this issue—which is entirely different from that of the current paper—please see Aumann, Hart, and Perry (1997).

6. DISCUSSION

This section elaborates on a number of different matters.

(a) *Decision Points*. Part of the problem in P & R's analysis is in their interpretation of information sets. Recall that the extensive form describes the way a game is *played*. The play proceeds from one node to the next, as each player is called upon to make a choice whenever a decision node of his is reached. Of course, when asked to make a choice, he possesses certain information. This is accurately described by information sets: two decision nodes where a player's information is the same belong to the same information set. But decisions are made at *nodes*, *not* at information sets (cf. the first observation of Section 2).

In games with perfect recall, a given information set can be reached only once—at a single node—in any one play of the game. Therefore, there is no harm in identifying decision points with information sets in such games, though even there the decision point is basically a node. But in games with absent-mindedness, when an information set may be visited more than once, it is simply incorrect to identify decision points with information sets.

(b) *Control*. At each intersection, the driver “expects” that he will do the same at the other intersection. He *expects* it, and maximizes given that expectation, and the maximizing behavior *turns out* the same as the expectation; that is precisely action-optimality. But *expecting* is not *determining*. He cannot, in fact, determine it—he cannot *control* what happens at the other intersection.

In their Section 7, P & R discuss what they call the “modified multi-selves approach,” which is the same as our notion of action-optimality. They write that this approach “*assumes* that a decision maker, upon reaching an information set, takes his actions to be immutable at future occurrences of that information set At the other extreme one finds the opposite axiom

⁷We are grateful to the associate editor in charge of the current paper for raising this question.

for which only one self resides in the information set and expects that, were the information set to occur again, he would adopt whichever behavior rule he adopts now" (their italics).

Piccione and Rubinstein's "opposite extremes" are in fact identical. The key element is *control*. At the action stage—once the driver reaches an intersection—there is no way that he can control or even affect what he does at the other intersection. So from his viewpoint, here and now, his future action really is "immutable." But that does not contradict P & R's EXIT "opposite axiom" (the "one self" approach). As we said above, *expecting* to do the same at the other intersection does not imply that the driver can *determine* here what happens there.

While P & R's "opposite axiom" is correct as written, their formalization of it is inappropriate. This formalization, which leads to "EXIT with probability 5/9," is based on the incorrect assumption that at the first intersection *X* the driver can *control* what he does at the second intersection *Y*.

(c) *Consistent Analysis*. Conceivably, P & R could challenge the first observation in Section 2—that at the action stage, the driver can control only what he does at the current intersection. Perhaps, after all, by some unspecified psychic process, he *can* control also what he does at a subsequent intersection. But in that case, he will have exercised this control already at the first intersection, and anything that he may think he is doing at the second intersection has no real effect. This makes the analyses at the action and planning stages identical, and then surely there is no paradox.

Another possibility is that the mysterious psychic process *affects* the decision at the subsequent intersection—say with some probability—but does not fully *determine* it. To analyze this possibility one would have to spell out just how the process works, and take it into account in the planning stage.

One could also consider a model in which the driver gets at most one chance to change his plan—at the first or second intersection, but not at both. In that case, too, he should take this into account in the planning stage, and again no paradox results.

In brief: One may consider various different scenarios. Whatever its specifications are, the precise scenario must be taken into account at the planning as well as at the action stage. The analyses at the action and planning stages must be consistent—they must analyze the same scenario.

(d) *The Pure Case*. Piccione and Rubinstein (P & R) probably agree that the pure case is not particularly interesting. Be that as it may, it turns out that in that case there is no action-optimal decision. If the driver EXITS at the other intersection, he should CONTINUE here, and if he CONTINUES at

the other intersection, he should EXIT here.⁸ (Formally, $h(p, 1)$ is maximized at $p = 0$, and $h(p, 0)$ at $p = 1$.)

Even though this is a one-person decision problem, randomized behavior is necessary at the action stage, because there are two independent decisions there. So how should the driver behave at the action stage? There is no clear answer. How do you play “Matching Pennies” when limited to pure strategies?

(e) *Tying Knots*. There is one particular scenario in Figure 1 that deserves further attention. Assume that the driver has a handkerchief in his pocket. Whenever he goes through an intersection, he ties a knot in the handkerchief, if there was no knot; or he unties the knot, if there was one. At the beginning (i.e., at START), it is equally probable that the handkerchief had or did not have a knot. The driver—absent-minded as he is—does not remember which was the case.

Thus, at each one of the two intersections, the probability of having a knot in the handkerchief is $1/2$. Therefore the driver does not learn anything about the intersection from the fact that there is or there is not a knot. The condition that “he does not know at which intersection he currently is” is satisfied.

However, the handkerchief allows the driver to use the following strategy: “EXIT if there is a knot, CONTINUE if there is not.” This is clearly better than anything he can do without the handkerchief (it yields a payoff of 2). The handkerchief has made it possible to separate the intersections without identifying them. It serves as an external correlation device. Of course, other things could be used—like sunspots, policemen, and so on (for instance, assume the single policeman in town chooses at random at which intersection to be).

In all these cases, the driver can behave differently at the two intersections. But then he should take this into account at the planning stage as well—and again there is no paradox (see Subsection (c) above).

⁸This may sound like P & R’s argument in the pure case, but it isn’t. Given the planning-optimal decision—which in the pure case is CONTINUE—P & R claim that at the action stage, the driver should switch to EXIT. They base this on $1/2 - 1/2$ probabilities of being at the two intersections; these probabilities are derived from the assumption that the driver indeed CONTINUES. But if he decides to EXIT, this assumption makes no sense: the probabilities cannot be computed as if he had chosen CONTINUE. As in Section 3, when calculating the expected payoff of switching plans, you cannot use probabilities as if you had not switched.

In contrast, we say that if the driver EXITS at the *other* intersection, he should CONTINUE here, and if he CONTINUES at the *other* intersection, he should EXIT here. That is a different kettle of fish altogether.

7. CONCLUSION

At the action stage, the driver must assume that the other decision is fixed at the action-optimal value. This is consistent with the optimal choice at the planning stage. Thus, the example of the absent-minded driver displays no dynamic inconsistency.

APPENDIX

We provide here the precise derivation of the function $h(p, q)$ of Section 4. Recall that p and q denote the probabilities of CONTINUE at the current and at the other intersection, respectively, and $h(p, q)$ denotes the resulting expected payoff at the current intersection. For now, think of p and q as fixed.

Define two random variables, \mathbf{Z} and \mathbf{t} . \mathbf{Z} is the end-node that is eventually reached, and \mathbf{t} is the current time. Thus \mathbf{Z} takes the values A , B , and C ; as for \mathbf{t} , we are only interested in two values, say $\mathbf{t} = 1$, which is the time when X is visited, and $\mathbf{t} = 2$, which is the time when Y is visited (if CONTINUE is chosen at X).

Without any information, it is equally probable that the current time \mathbf{t} is 1 or 2:

$$P(\mathbf{t} = 1) = P(\mathbf{t} = 2) = 1/2;$$

otherwise, the two decision points would be distinguishable. This holds for the total probability (not conditional on one end-node or another).

From the definition of p and q , we have:

$$\begin{aligned} P(\mathbf{Z} = A | \mathbf{t} = 1) &= 1 - p, & P(\mathbf{Z} = A | \mathbf{t} = 2) &= 1 - q, \\ P(\mathbf{Z} = B | \mathbf{t} = 1) &= p(1 - q), & P(\mathbf{Z} = B | \mathbf{t} = 2) &= q(1 - p), \\ P(\mathbf{Z} = C | \mathbf{t} = 1) &= pq, & P(\mathbf{Z} = C | \mathbf{t} = 2) &= qp. \end{aligned}$$

Putting it all together yields

$$\begin{aligned} P(\mathbf{Z} = A \text{ and } \mathbf{t} = 1) &= (1 - p)/2, \\ P(\mathbf{Z} = A \text{ and } \mathbf{t} = 2) &= (1 - q)/2, \\ P(\mathbf{Z} = B \text{ and } \mathbf{t} = 1) &= p(1 - q)/2, \\ P(\mathbf{Z} = B \text{ and } \mathbf{t} = 2) &= q(1 - p)/2, \\ P(\mathbf{Z} = C \text{ and } \mathbf{t} = 1) &= pq/2, \\ P(\mathbf{Z} = C \text{ and } \mathbf{t} = 2) &= qp/2. \end{aligned}$$

The ‘‘current intersection’’ \mathbf{N} is defined as the intersection, if any, visited at the current time \mathbf{t} . If $\mathbf{t} = 1$, it is necessarily X ; if $\mathbf{t} = 2$, it is Y if $\mathbf{Z} = B$

or $\mathbf{Z} = C$, and “none” if $\mathbf{Z} = A$ (we write this as $\mathbf{N} = \emptyset$). Thus we obtain $P(\mathbf{N} = X) = P(\mathbf{t} = 1) = 1/2$; $P(\mathbf{N} = Y) = P(\mathbf{t} = 2 \text{ and } \mathbf{Z} \in \{B, C\}) = q(1 - p)/2 + qp/2 = q/2$; and $P(\mathbf{N} = \emptyset) = P(\mathbf{t} = 2 \text{ and } \mathbf{Z} = A) = (1 - q)/2$.

The expected payoff $h(p, q)$ at the current intersection can thus be written as

$$h(p, q) = P(\mathbf{N} = X | \mathbf{N} \in \{X, Y\}) E(u(\mathbf{Z}) | \mathbf{N} = X) \\ + P(\mathbf{N} = Y | \mathbf{N} \in \{X, Y\}) E(u(\mathbf{Z}) | \mathbf{N} = Y),$$

where $u(\mathbf{Z})$ is the payoff at the end-node \mathbf{Z} . Therefore

$$h(p, q) = \frac{1/2}{1/2 + q/2} [(1 - p) \cdot u(A) + p(1 - q) \cdot u(B) + pq \cdot u(C)] \\ + \frac{q/2}{1/2 + q/2} [(1 - p) \cdot u(B) + p \cdot u(C)].$$

Thus, conditional on currently being at an intersection, the probability is $1/(1 + q)$ that it is X , and $q/(1 + q)$ that it is Y . These are the “consistent beliefs” of P & R. (Note that the beliefs about the identity of the current intersection depend only on the behavior at the *other* intersection—not at the current one, where nothing has yet been done. Hence these probabilities are a function of q and not of p .)

Next, let x and y denote the probabilities of CONTINUE at X and Y , respectively. The expected payoff (at START) is then

$$f(x, y) := (1 - x) \cdot u(A) + x(1 - y) \cdot u(B) + xy \cdot u(C).$$

A behavior \bar{p} is planning-optimal if it maximizes $f(p, p)$ over p . To compare planning-optimality with action-optimality, note that

$$\left(\frac{1}{2} + \frac{q}{2}\right) \cdot h(p, q) + \left(\frac{1}{2} - \frac{q}{2}\right) \cdot u(A) = \frac{1}{2}f(p, q) + \frac{1}{2}f(q, p).$$

Denote this expression by $g(p, q)$; the right side may be interpreted as the expected payoff, evaluated at START, of choosing p at one intersection and q at the other, but without knowing which is which. Now p^* is action-optimal if it maximizes $h(p, p^*)$ over p , the second argument being fixed at p^* . Equivalently, since $h(p, p^*)$ is, for fixed p^* , a positive linear transformation of $g(p, p^*)$ (see above), it follows that p^* is action-optimal if and only if p^* maximizes $g(p, p^*) = [f(p, p^*) + f(p^*, p)]/2$ over p . This implies that the randomized planning-optimal decision \bar{p} is action-optimal. Indeed, the first-order necessary conditions of the two problems are identical; they are moreover sufficient for action-optimality, where the function to be maximized is linear.⁹

⁹This argument is general; it proves Proposition 3 of P & R.

REFERENCES

- Aumann, R. J., Hart, S., and Perry, M. (1997), "The Forgetful Passenger," *Games and Econ. Behav.* **20**, 117–120.
- Piccione, M., and Rubinstein, A., (1997). "On the Interpretation of Decision Problems with Imperfect Recall," *Games and Econ. Behav.* **20**, 3–24.