The Absent - Minded Driver[1]

by

ROBERT J. AUMANN, SERGIU HART AND MOTTY PERRY

The Hebrew University of Jerusalem[2]

January 1996

Hopefully without ... what was it?
Ah! ... absent-mindedness ...

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 & Rubinstein [1995; P&R in the sequel], 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

[1] 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.
[2] Center for Rationality and Interactive Decision Theory, The Hebrew University of Jerusalem, Feldman building, Givat-Ram, 91904 Jerusalem, Israel; fax: +972-2-6513681; e-mail: ratio@vms.huji.ac.il.

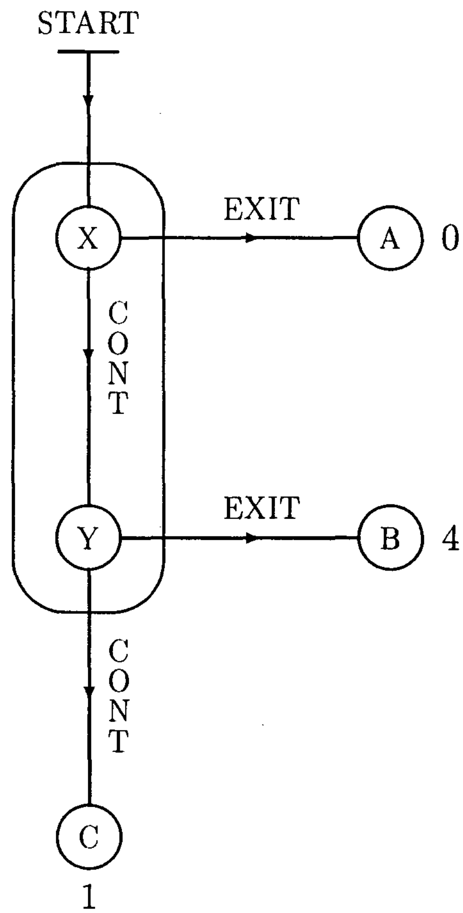

Figure 1. The absent-minded driver problem

provocative and worth having, P&R's analysis is unsound. A careful analysis reveals that while the considerations at the planning and action stages do differ, there is no paradox or inconsistency.

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. The correct analysis is provided in Section 4 -- where we formally define the concept of *action-optimality* -- and in Section 5. In Section 6 we study action-optimality in a slightly more general setup -- with some interesting and unexpected conclusions. We conclude with a detailed discussion of various issues in Section 7.

2. Fundamentals

At the **planning** stage, the decision problem is straightforward. In the example, the optimal randomized decision is "CONTINUE with probability 2/3 and EXIT with probability 1/3", and the optimal pure decision is "CONTINUE"[3]. We call these the *planning-optimal* decisions.

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[4].

> • **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.

---

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

[4] One 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. P&R have yet to adduce an explicit scenario that *does* display a paradox. See Section 7(c) for more on this issue.

• **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

Assume first, as in P&R, that the driver is limited to pure strategies. Consider the action stage. The driver finds himself at an intersection; he does not know which. P&R argue that X and Y are equally likely -- have probability 1/2 each -- and thus the expected payoff from EXIT is $0 \cdot (1/2) + 4 \cdot (1/2) = 2$, which is better than the payoff of 1 from CONTINUE. Therefore the driver should EXIT. But where do these (1/2, 1/2) probabilities come from? Clearly, from the assumption that the driver chooses CONTINUE. But if he decides to EXIT, this assumption makes no sense: the probabilities cannot be computed as if he had chosen CONTINUE! P&R's fallacy is that the driver decides to do something, based on the belief that he is doing something else -- a violation of the second observation.

Next, consider the case where randomization is allowed. The expected payoff at the action stage is

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

where α is the probability of being currently at X, and p and q are the respective probabilities of CONTINUE at the current intersection and at the "other" intersection. P&R maximize $H(p,p,\alpha)$ 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 it is unjustified to treat α as if it were independent of that probability.

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^*.

(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 is[5] $\alpha = 1/(1+q)$. 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 point[6] of the (set-valued) mapping $q \longrightarrow \arg\max_p h(p,q)$.

Applying this analysis to the example, we see that the 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 $h(p,2/3) = (3/5)\cdot[0\cdot(1-p) + 4\cdot p\cdot(1/3) + 1\cdot p\cdot(2/3)] + (2/5)\cdot[4\cdot(1-p) + 1\cdot p]$, which equals 8/5 (for all p). So $p = 2/3$ maximizes it; thus $p^* = 2/3$ is action-optimal.

---

[5] For a derivation of this probability, see the Appendix. Informally, since the driver always goes through X, but only a proportion q of the time through Y, the ratio of the probabilities is 1 to q, thus they are $1/(1+q)$ and $q/(1+q)$, respectively. These are the "consistent beliefs" of P&R; in the Appendix we show that these are the only possible probabilities.

[6] To restate this formally in familiar game-theoretic terms, (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".

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,

$$h(p,q) = \frac{1}{1+q}\cdot\left[0\cdot(1-p) + 4\cdot p(1-q) + 1\cdot pq\right] + \frac{q}{1+q}\cdot\left[4\cdot(1-p) + 1\cdot p\right]$$
$$= \frac{(4-6q)\cdot p + 4q}{1+q} .$$

Given q, the maximizing p is therefore 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 multi-selves 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 7 below for further discussion.

5. The Pure Case

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 decision were to EXIT, then, given that EXIT is chosen at the other intersection, it is better to CONTINUE at the current intersection (indeed, the current intersection cannot be Y -- since that happens only if CONTINUE was chosen at the other intersection; but then CONTINUE now leads to a higher payoff than EXIT). Similarly, if the decision were to CONTINUE, then it is better to EXIT now. (Formally, the maximum of $h(p,1)$ is attained at $p = 0$, and that of $h(p,0)$ at $p = 1$.)

Even though this is a one-person decision problem, the fact that, from the viewpoint of the action stage, there are two independent decisions, makes randomized behavior necessary. So how should the driver behave at the action stage? There is no clear answer. How does one play "Matching Pennies" when limited to pure strategies?

6. A More Challenging Example

In Section 4, we saw that, in the specific example of P&R, randomized action-optimality and planning-optimality coincide. In general, any planning-optimal randomized decision is also action-optimal -- see Proposition 3 of P&R (and also the end of the Appendix below). However, there may be additional action-optimal choices. For instance -- see Figure 2 -- change the payoffs to be 1 at A, 0 at B and 2 at C. 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).

The question now arises: How do the payoffs of the various action-optimal decisions compare? Of course, when computed at START, those that are also planning-optimal yield the maximal payoff. But this need not be so for the payoff computed from the viewpoint of *the current intersection*, i.e. $h(p^*,p^*)$. Such an example requires three intersections; take the payoffs to be 7 at the first EXIT, 0 at the second EXIT, 22 at the third EXIT, and 2 if always CONTINUE (see Figure 3). It may be checked that:

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

(ii) There are three action-optimal decisions, namely $p_1^* = \bar{p} = 0$, $p_2^* = 7/30$ 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 $h(p_1^*,p_1^*) \equiv h(\bar{p},\bar{p})$.

The reader may ask, since the choice is among three possibilities yielding 7, $\approx$6.37 or $\approx$7.14, why doesn't the driver choose the action with the highest yield, namely p_3^*? The answer, of course, is that at the action stage,

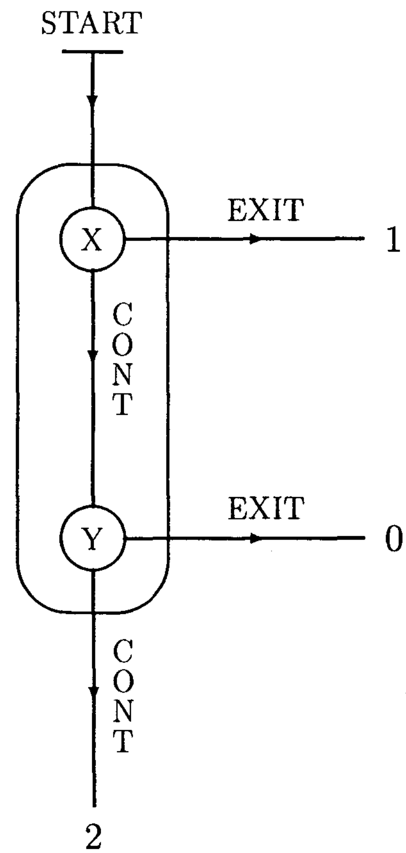

Figure 2. Multiple action-optimal decisions

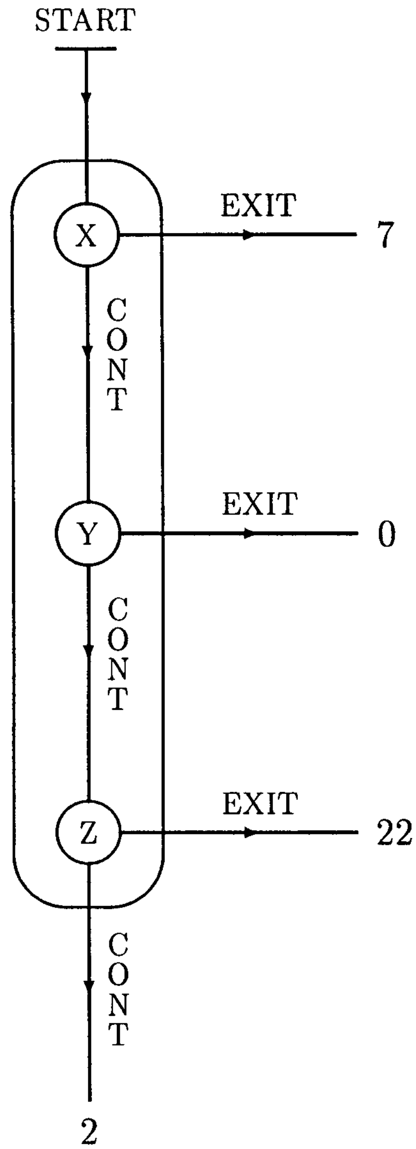

Figure 3. A more challenging example

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 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, 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^*, then 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; the result should be one of the p_i^*.

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?

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 only because 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 Figure 4. The car is automatic and exits with probability 1/2 at each intersection. The decision maker is 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 refer to the second option as "clear-headedness".

Like 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 to CONTINUE is increased to 4/7 (Figure 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

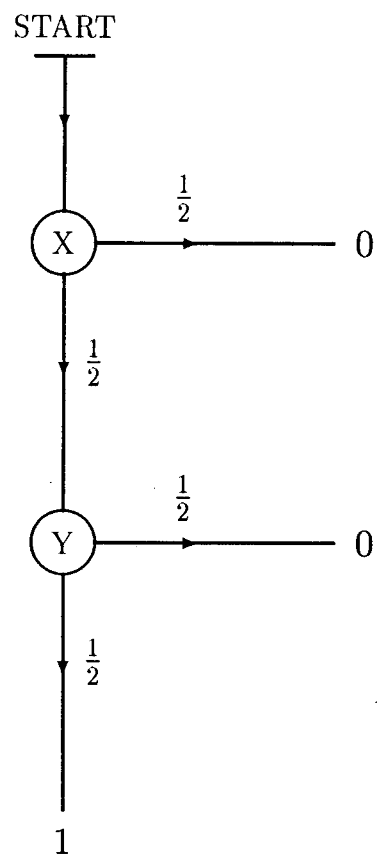

Figure 4. An automatic car

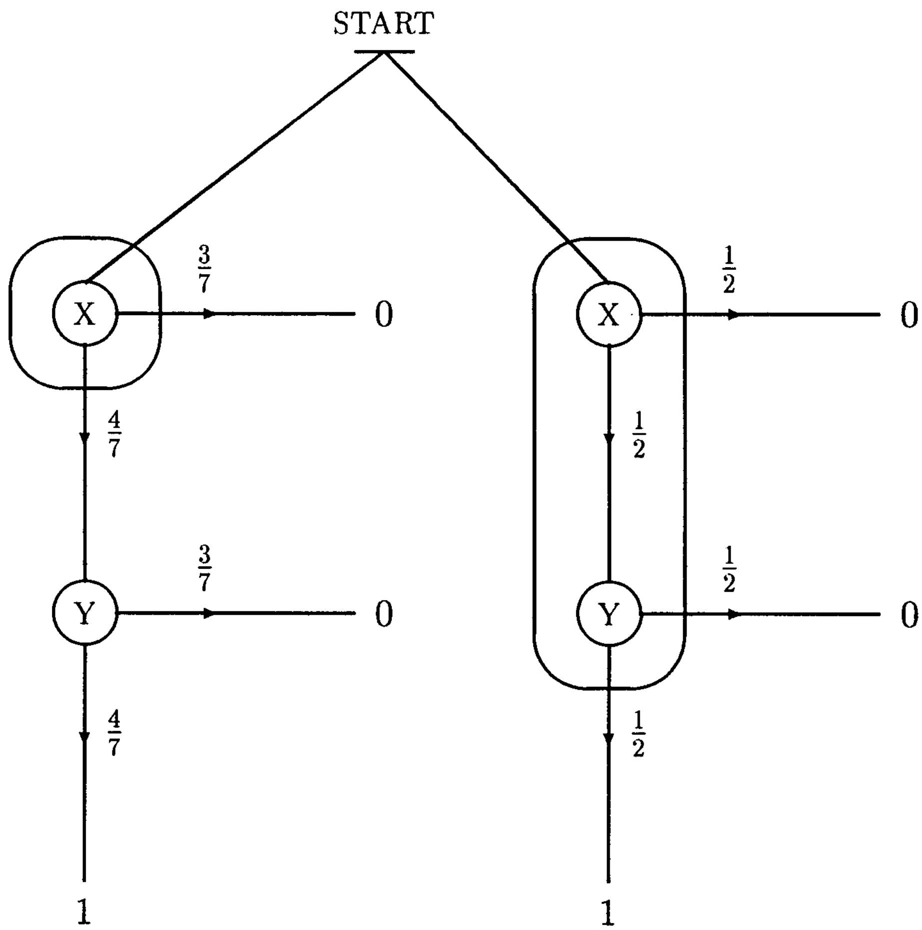

Figure 5. Clear-headed or absent-minded?

decision maker to wish he were absent-minded -- it would be sticking his head in the sand!

7. Discussion

(a) Decision points. P&R's mistake is in the way they interpret 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. 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, 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." (P&R, page 21; their italics).

P&R'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 "opposite axiom". At each intersection, the driver indeed "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.

While P&R's "opposite axiom" is correct as written, their formalization of it is not. 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 this effect works, and take it into account in the planning stage. P&R have not done this.

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 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 the action stage. The analyses at the action and planning stages must be consistent -- they must analyze the same scenario. P&R's don't; ours do.

(d) Tying knots. There is one particular scenario 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 isn't 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 isn't". 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 between 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 has the possibility to 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).

8. 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. Moreover, one must be consistent and analyze the same scenario at both stages -- planning and action.

APPENDIX

We provide here the precise derivation of the function h(p,q) of Section 4. Recall that p and q denote the probabilities to 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, **Z** and **t**. **Z** is the end-node that is eventually reached, and **t** is the current time. Thus **Z** takes the values A, B and C; as for **t**, we are only interested in two values, say **t** = 1, which is the time when X is visited, and **t** = 2, which is the time when Y is visited (if CONTINUE is chosen at X).

From the definition of p and q, we have:

$$P(\mathbf{Z} = A|\mathbf{t} = 1) = 1-p, \qquad P(\mathbf{Z} = A|\mathbf{t} = 2) = 1-q,$$

$$P(\mathbf{Z} = B|\mathbf{t} = 1) = p(1-q), \qquad P(\mathbf{Z} = B|\mathbf{t} = 2) = q(1-p),$$

$$P(\mathbf{Z} = C|\mathbf{t} = 1) = pq, \qquad P(\mathbf{Z} = C|\mathbf{t} = 2) = qp.$$

In addition, it is equally probable that the current time **t** is 1 or 2:

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

Otherwise, the two decision points would not be indistinguishable[7]. Of course, this is true when we consider the total probability, not the conditional probabilities given one end-node or another.

Putting it all together yields:

$P(\mathbf{Z} = A \text{ and } \mathbf{t} = 1) = (1-p)/2,$ $\quad P(\mathbf{Z} = A \text{ and } \mathbf{t} = 2) = (1-q)/2,$

$P(\mathbf{Z} = B \text{ and } \mathbf{t} = 1) = p(1-q)/2,$ $\quad P(\mathbf{Z} = B \text{ and } \mathbf{t} = 2) = q(1-p)/2,$

$P(\mathbf{Z} = C \text{ and } \mathbf{t} = 1) = pq/2,$ $\quad P(\mathbf{Z} = C \text{ and } \mathbf{t} = 2) = qp/2.$

The "current intersection" N is defined as the intersection, if any, visited at the current time **t**. If **t** = 1, it is necessarily X; if **t** = 2, then it is Y if **Z** = B or **Z** = C, and "none" if **Z** = A (we write this as **N** = ø). Thus we obtain $P(N = X) = P(\mathbf{t} = 1) = 1/2$; $P(N = Y) = P(\mathbf{t} = 2 \text{ and } \mathbf{Z} \in \{B, C\}) = q(1-p)/2 + qp/2 = q/2$; and $P(N = \emptyset) = P(\mathbf{t} = 2 \text{ and } \mathbf{Z} = A) = (1-q)/2$.

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

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

where u(Z) is the payoff at the end-node Z. Thus

[7] A more formal argument is as follows: Denote the driver's two decision points by I and J (one is at X and the other at Y -- but it is not known which is which). In each he possesses certain knowledge and certain probabilities. Let φ be the probability $P_I(I=X)$ that he ascribes at I to the event that I is the first intersection (i.e., X). Since I and J are indistinguishable, the probability $P_J(J=X)$ that he ascribes at J to J being X, is also φ. Since I=X and J=X are complementary events, it follows that $P_J(I=X) = 1-\varphi$. Since we are talking about one and the same driver, we can assume that the driver at I and the driver at J have common priors, and that $P_I(I=X)$ and $P_J(I=X)$ are commonly known between the driver at I and the driver at J. Therefore by the Agreement Theorem (Aumann, 1976), we have $\varphi = P_I(I=X) = P_J(I=X) = 1-\varphi$. So $\varphi = 1/2$.

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

$$\frac{q/2}{1/2 + q/2}\Big[u(B)\cdot(1-p) + u(C)\cdot p\Big].$$

Note that we have obtained that, given that one is currently at an intersection, the probability that it is X equals 1/(1+q), and that it is Y equals q/(1+q). These are the "consistent beliefs" of P&R (indeed, which is the current intersection depends on the behavior only at the *other* intersection, and not at the current one -- where nothing has been yet done; therefore these probabilities are a function of q and not of p).

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

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

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

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

Let g(p,q) denote this expression; the right-hand 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 is a maximizer of h(p,p*) over p, the second argument being fixed at p*. Equivalently, since h(p,p*), for fixed p*, is just a positive linear transformation of g(p,p*) (see above), we have: p* is action-optimal if and only if p* is a maximizer of g(p,p*) = [f(p,p*) + f(p*,p)]/2 over p. From this it immediately follows 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[8].

[8] This argument is general; it proves Proposition 3 of P&R.

REFERENCES

Aumann, R. J. [1976], Agreeing to Disagree, *Annals of Statistics* 4, 1236-1239.

Piccione, M. and A. Rubinstein [1995], "On the Interpretation of Decision Problems with Imperfect Recall", WP 24-94 (revised September 1995), Tel-Aviv University; *Games and Economic Behavior*, forthcoming.