

Rational Updating at the Crossroads ¹

EPICENTER Working Paper No. 23 (2021)



Silvia Milano* and Andrés Perea[†]

ABSTRACT. In this paper we explore the absentminded driver problem using two different scenarios. In the first scenario we assume that the driver is capable of reasoning about his degree of absentmindedness before he hits the highway. This leads to a Savage-style model where the states are mutually exclusive and the act-state independence is in place. In the second we employ centred possibilities, by modelling the states (i.e. the events about which the driver is uncertain) as the possible final destinations indexed by a time period. The optimal probability we find for continuing at an exit is different from almost all papers in the literature. In this scenario, act-state independence is still violated, but states are mutually exclusive and the driver arrives at his optimal choice probability via Bayesian updating. We show that our solution is the only one guaranteeing immunity from sure loss via a Dutch strategy, and that – despite initial appearances – it is time consistent.

1. INTRODUCTION

Ever since its introduction by Piccione and Rubinstein (1997), the *absent-minded driver problem* has provoked lively discussions and controversy among game theorists, decision theorists, logicians, philosophers, computer scientists and economists. The problem seems thought provoking for several reasons.

First, it stands apart from traditional one-person decision problems under uncertainty, as in most of its existing formulations the independence between acts and states is violated, and the states are not mutually exclusive. Indeed, upon reaching an exit on the highway, the absentminded driver is uncertain whether this is the first or second exit, but his act of whether or not to take the exit influences the probability of reaching the second exit. Hence, if the states merely represent the two exits, then the act chosen influences the likelihood of the two states. Moreover, the states would also not be mutually exclusive in this case, because reaching the second exit implies that the first exit must have been reached previously. This poses a problem for how the driver must update his beliefs upon seeing an exit.

¹We would like to thank three anonymous referees at the TARK XVIII conference for their useful comments.

*Future of Humanity Institute, University of Oxford.

E-mail: silvia.milano@philosophy.ox.ac.uk Web: <https://www.silvia-milano.com/>

[†]Department of Quantitative Economics, Maastricht University.

E-mail: a.perea@maastrichtuniversity.nl Web: <https://www.epicenter.name/Perea/>

Second, there seems to be a tension between planning the decision ex-ante, and implementing the decision once the driver is on the highway. Most of the existing papers find that the ex-ante optimal probability of leaving the highway is no longer optimal if the driver really sees the exit, and updates his beliefs accordingly. This presents a serious problem.

In this paper we present two possible resolutions to the problem. In the first, we assume that the driver can reason about his degree of absentmindedness before making a decision. More precisely, the driver contemplates the possibilities that he would be absentminded only at the first exit, only at the second exit, at both exits or at no exit. These four states are mutually exclusive, and restore the act-state independence. This allows us to undertake a traditional Savage-style analysis (Savage (1972)). In particular, we find that the tension between ex-ante planning and optimal choice on the highway disappears.

In our second scenario, we assume that the states do not merely correspond to the two exits, but represent the three possible final destinations indexed by time. Assuming there are two time periods (the time when the first exit would appear, and the time when the second exit would appear), this results in six possible *centred states*. Moreover, the six states would be mutually exclusive (but act-state independence would still be violated). If the driver sees an exit, then he must rule out the centred state where the first exit has been taken and the time period is 2. We find that, upon revising the belief by Bayesian updating, the optimal probability of continuing when seeing an exit is approximately 0.53. This is different from most papers in the literature which point at an optimal continuation probability of $\frac{2}{3}$. We also show that adopting our solution is the only way in which the driver can be immune to diachronic Dutch strategies.

The outline of the paper is as follows. In Section 2 we introduce the absentminded driver problem, followed by a discussion of some of the relevant literature. In Sections 3 and 4 we analyse the absentminded driver problem using the two different scenarios described above. In Section 5, we analyse when the driver would be immune to Dutch strategies, and then consider the implications of our solution for time consistency in Section 6. Section 7 concludes.

2. THE ABSENTMINDED DRIVER PROBLEM

Consider a person who must drive home after visiting a friend. On the highway there are two exits, and to reach his house he must leave the highway at the second exit. However, since he is absentminded, he cannot remember at the second exit whether he saw the first exit or not. As both exits look alike, he does not know at which of the two exits he is when he sees one. If he leaves the highway at the first exit he will enter a scary road – something that will give him nightmares for the rest of his life. If, on the other hand, he continues at the second exit then he arrives at a motel. He clearly prefers to sleep at home rather than spending the night at the motel. This situation can be visualised by Figure 2.1. The question is: What should the driver do when he sees an

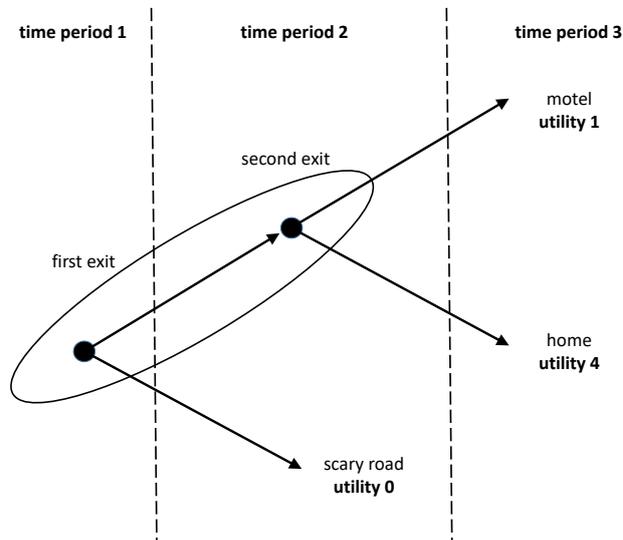


Figure 2.1: The absentminded driver problem

exit? Continue or leave the highway, or perhaps even randomise over these two options? This is known as the *paradox of the absentminded driver*, or simply the *absentminded driver problem*. See Piccione and Rubinstein (1997).

We now turn to review some of the existing approaches to the absentminded driver problem. Subsequently, we point at some potential problems with these approaches.

2.1. THE TIME CONSISTENCY PROBLEM

Piccione and Rubinstein (1997) point out that there seems to be a tension between optimal planning and time consistency in the absentminded driver problem. Suppose the driver must plan what to do before he leaves the house of his friend. Since he cannot make his choice conditional on reaching the first or second exit, a strategy would simply be: continue if you see an exit, leave the highway if you see an exit, or possibly a randomisation over these two possibilities. Let us first focus on the case without randomisation. Then, the best strategy from an ex-ante perspective would be to continue if you see an exit. Now, suppose the driver finds himself on the highway and sees an exit. Given his strategy to always continue, it seems reasonable for him to believe that he is at the first exit with probability 0.5. But then, the expected utility of leaving the highway would be 2, which is more than what he would get by continuing. Hence, the driver would be tempted to change his plan once he finds himself on the highway. In other words, the optimal plan from the ex-ante perspective is not time consistent.

This tension between optimal planning and time consistency will persist if we

allow for randomisation. Suppose the driver plans to use the same randomisation device every time he sees an exit, which induces him to continue with probability π and to leave with probability $1 - \pi$ at both exits. Then, the ex-ante expected utility from this plan would be $\pi^2 \cdot 1 + \pi(1 - \pi) \cdot 4 + (1 - \pi) \cdot 0$, which is maximised when $\pi = \frac{2}{3}$. Suppose again that the driver finds himself on the highway and sees an exit, and let β be the probability that he assigns to being at the first exit. Given his strategy, it seems reasonable for him to believe that the probability of being at the second exit is $\beta \cdot \frac{2}{3}$, which means that $\beta = 0.6$. Given these beliefs, however, the expected utility of leaving the highway is higher than that of continuing, and hence the driver would again be tempted to change his plan to leaving the highway with probability 1. That is, also the probabilistic ex-ante optimal plan is not time consistent.

2.2. THE MODIFIED MULTISELF CONSISTENCY APPROACH

As a possible resolution to this tension, Piccione and Rubinstein focus, at the end of the paper, on a different type of time consistency, called “modified multiseLF consistency”. The difference with the usual time consistency is that the driver is now modelled as a system with two agents: one who takes the decision at the first exit, and one who takes the decision at the second exit. Of course, “in equilibrium”, both agents must reach the same decision, as the driver cannot distinguish between the two exits. But the optimality criterion is different now: Instead of requiring that the driver cannot do better, given his beliefs, by uniformly changing his decision at both exits, it is now required that the first agent cannot do better, given his beliefs and *given the fixed randomisation of the second agent*. Strategies that satisfy this criterion are called “modified multiseLF consistent”.

It may be verified that the ex-ante optimal plan, to continue at each exit with probability $\frac{2}{3}$, is modified multiseLF consistent. Indeed, given the belief of being at the first exit with probability 0.6, and given the belief that the second agent will continue with probability $\frac{2}{3}$, the expected utilities for the first agent of continuing and leaving are both 1.6, and hence it remains optimal for the first agent to randomise.

Whereas Piccione and Rubinstein (1997) do not state a clear preference for either their standard form of time consistency or modified multiseLF consistency, Aumann et al. (1997) claim that the latter notion of consistency is the more natural one. In their paper this type of consistency is called “action-optimality”. They argue that if modified multiseLF consistency is adopted as the *only* reasonable time consistency criterion, then the paradox in the absent-minded driver problem disappears. Indeed, as we have seen above, the unique ex-ante optimal plan is also modified multiseLF consistent.

Also Gilboa’s (1997) analysis eventually leads to the criterion of modified multiseLF consistency, but via a different route. Gilboa transforms the decision tree from Figure 2.1 into a two-player game that satisfies perfect recall, by having nature determine at random which of the two agents will act at the first exit.

The two agents act as the two players in the game, with identical utilities at the end. In particular, an agent believes that he will never have to act twice in a row, removing absentmindedness – and thereby imperfect recall – from the game. The symmetric Nash equilibria of this modified two-player game correspond precisely to the strategies in the original decision tree that are modified multiselected.

2.3. OTHER APPROACHES

As we have seen above, the ex-ante optimal plan to continue with probability $\frac{2}{3}$ is no longer optimal if the driver finds himself on the highway, at least if we do not view the driver as a system with two separate agents. Although we will argue otherwise later (see Section 6 below), this may appear paradoxical, as it may seem that the driver does not learn anything substantially new if he finds himself on the highway. Indeed, the driver already knew at the ex-ante stage that he would eventually hit the highway. So why should the driver’s doxastic state at the highway lead to a different choice than at the ex-ante stage?

Rabinowicz (2003), however, argues that the driver’s doxastic state at the highway *is* fundamentally different from the ex-ante stage, because the driver at the highway realises that he may no longer be able to fully implement the strategy he planned at the ex-ante stage. To be more precise, a strategy describes the (deterministic or probabilistic) choice the driver plans to make *each time* he sees an exit. However, when he is at the highway and sees an exit, he realises this may well be the second exit already, which would mean that he is no longer able to fully implement the strategy at both exits. As, according to Rabinowicz, the driver’s mental state at the ex-ante stage is fundamentally different from when he finds himself on the highway, it is argued that there is no real paradox here. As will become clear later on, we argue for a similarly ‘deflationist’ resolution of the absent-minded problem (see Sections 4-6), but our approach differs from Rabinowicz as it does not involve modelling the driver’s internal mental states.

Halpern and Pass (2016) propose the new concept of “sequential equilibrium” for dynamic games with imperfect recall. This includes the absentminded driver problem as a special case. An important ingredient of their concept is that at information sets containing consecutive nodes – like the one in the absentminded driver problem – the player’s belief should only assign positive probability to nodes that come first. In the absentminded driver example, this would mean that the driver’s belief should assign probability 1 to the first exit. Clearly then, the decision problem for the driver at his information set would be the same as at the ex-ante stage, removing the tension between optimal planning and time consistency from the outset.

Finally, Schwarz (2015) investigates the absentminded driver problem by considering three different states: being at the first exit and believing that you would continue at the second exit, being at the first exit and believing that you would leave at the second exit, and being at the second exit. Hence, the

driver’s counterfactual action at the second exit is included in the description of a state. Schwarz analyses the problem by using two different approaches: a causal decision theory approach (Savage, 1972; Lewis, 1981) and an evidential decision theory approach (Jeffrey, 1990). The main difference is that in the first approach, the driver holds a fixed, subjective probabilistic belief about the three states which he then uses to determine what to do. In the evidential decision theory approach, in turn, the driver bases his belief about the states on the action he is planning to take.

Schwarz starts by investigating the scenario where no randomisation is allowed. The causal decision theory approach then leads to an “unstable” decision problem where the attractiveness of a choice is inversely related to the probability assigned to that choice. In particular, there is a cut-off probability for continuing such that continuing is optimal precisely when the probability with which the driver believes to continue is below the cut-off. It is argued that some dynamics of rational deliberation à la Skyrms (1990) may lead the driver to an equilibrium state of indecision in which the probability with which he believes to continue is equal to the cut-off probability. At such a state, the driver would thus be indifferent between continuing and leaving the highway at an exit. Depending on whether a certain uniformity condition is assumed or not, the cut-off probability for continuing will be either $\frac{2}{3}$ or $\frac{4}{5}$. The evidential decision theory approach, in contrast, leads to an unambiguous recommendation, which is to continue at the exit.

Subsequently, the analysis is extended to the setting where the driver is allowed to randomise between the two options “continue” and “leave”. This is simulated by letting the driver use a biased coin at both exits. The set of states must then be enlarged, as a state must also specify which bias the driver will be using at the second exit if he is currently at the first exit. The causal decision theoretic analysis becomes analogous to the modified multisection approach in Piccione and Rubinstein (1997), leading to an optimal coin with bias $\frac{2}{3}$ towards continuing. The evidential decision theory approach, however, comes to a different conclusion. The reason is that within this approach, the driver must believe that whatever bias he chooses, he will always use that same bias at both exits. Schwarz then shows that the optimal bias depends on the way in which the driver calculates the conditional probabilities relevant to his decision, in a way that mirrors the difference between so-called ‘thirder’ and ‘halfer’ positions in the Sleeping Beauty problem. While the solution that we will put forward in Section 4 below agrees with Schwarz’s combination of evidential decision theory plus the ‘thirder’ conditional probabilities, we arrive at this value through a different reasoning. In particular, in our model the states do not include counterfactual information, and our solution is not dependent on the choice of evidential decision theory.

2.4. DISCUSSION

Each of the approaches above, except Schwarz (2015) perhaps, may be viewed as an attempt to remove the tension between ex-ante optimality and time

consistency in the absentminded driver problem. In our view, however, there are some conceptual problems with these approaches.

Savage (1972) assumes, in his axiomatisation of expected utility, that the decision maker cannot influence the state of the world by the act he chooses. It is precisely this act-state independence that is being violated in each of the papers above, except Gilboa (1997).² Indeed, in each of these papers except Schwarz (2015), if the driver finds himself on the highway, then the relevant states of the world are the events “I am at the first exit” and “I am at the second exit”. Given the time periods reproduced in Figure 2.1, the first and second state can alternatively be described as “time period is 1” and “time period is 2 and I continued at time period 1”. Clearly, whether or not the second state is realised crucially depends on which act is chosen. That is, Savage’s act-state independence is violated. In Schwarz (2015), the relevant states are extended to “I am at the first exit and believe to continue at the second exit”, “I am at the first exit and believe to leave at the second exit”, and “I am at the second exit”. Also here, the act chosen will influence the state that is realised.

Another assumption in Savage’s framework is that the states of the world represent *mutually exclusive* events. However, in all of the papers above, except Gilboa (1997), the states are not mutually exclusive from an ex-ante perspective. Indeed, before leaving the friend’s house, the event of reaching the second exit is perfectly compatible with – in fact, can only happen after – reaching the first exit. This fact is explicitly acknowledged by Piccione and Rubinstein when discussing their notion of consistent beliefs: When deriving the consistent belief probabilities from a given strategy, they state that the total ex-ante probability of reaching the first or second exit may be larger than 1 (in fact, will be larger than 1), precisely because these two events are not mutually exclusive from an ex-ante perspective. In that sense, their notion of consistency is different from standard Bayesian updating, as this only applies to mutually exclusive events.

As an alternative way of updating beliefs, Piccione and Rubinstein introduce the notion of “*Z*-consistency”, and show by means of examples that it differs from consistency. They state that they are not sure which of these two notions of consistency is more appealing. This ambiguity as to which updating rule must be used is a direct consequence of the states of the world not being mutually exclusive from an ex-ante perspective.

The consistency notion used by Halpern and Pass (2016) is also different from the usual Bayesian updating, as they require the driver to assign probability 1 to the first exit, irrespective of his plan. However, we believe that this restriction goes against the essence of the absentminded driver problem, which is that the driver is inherently uncertain about the exit he sees, and takes this information to be relevant for what he should choose.

²Since Gilboa’s (1997) analysis splits the driver into two distinct decision makers, and each decision maker faces his own personal set of states which can only be affected by the choice of the other decision maker, act-state independence is still preserved in this paper.

	ω_0	ω_1	ω_2	ω_{12}
continue	motel (1)	motel (1)	motel (1)	motel (1)
leave	scary road (0)	home (4)	scary road (0)	motel (1)

Table 1: Reasoning about your degree of absentmindedness

Finally, we see some conceptual problems with the multiself approach that is being used in each of the papers above, except Halpern and Pass (2016). The main idea is that the driver is split into two separate agents, one choosing at the first exit and one at the second exit. In Gilboa (1997), these two agents are even explicitly modelled as two separate players. Moreover, if from a given plan an agent compares his planned choice with an alternative, counterfactual choice, then he assumes that the other agent will still stick to his planned choice. But how reasonable is this assumption? At the end, the two agents are time-slices of the same person, and so it is natural to think that the reasoning strategies and cognitive resources at their disposal are identical. If one arrives at a unique rational strategy, the other would also reach the same conclusion, and hence it seems likely that agent 1 could counterfactually change his choice along with agent 2.

3. REASONING ABOUT YOUR DEGREE OF ABSENTMINDEDNESS

We have seen that in each of the existing approaches in the literature to the absentminded driver problem, either act-state independence was violated, or the states were not mutually exclusive, or the driver was separated into two agents that choose independently. As a consequence, none of these approaches can be fully reconciled with the Savage (1972) framework.

In this section we offer an alternative approach that is completely in line with the Savage set-up. We do not explicitly incorporate time into our analysis, but rather allow the driver to reason about his degree of absentmindedness. This will enable us to provide a Savage-style model in which the events about which the driver is uncertain at the ex-ante stage are mutually exclusive, and where the driver's possible acts – to continue or to leave the highway at an exit – do not influence the state. Moreover, the driver is not separated into different agents.

More precisely, at the ex-ante stage the driver believes that he may either be absentminded (i) at no exit, (ii) at the first exit only, (iii) at the second exit only, or (iv) at both exits. If the driver is absentminded at an exit then, as a consequence, he will continue at this exit even if he planned to leave there, and he will forget about this exit for the remainder of the journey. The four mutually exclusive events listed above correspond to the states $\omega_0, \omega_1, \omega_2$ and ω_{12} , respectively. The decision problem at the ex-ante stage may therefore be summarised by Table 1.

Note that there is no dependence between acts and states in this model. Indeed,

histories	motel (A)	home (B)	scary road (C)
probabilities	π^2	$\pi - \pi^2$	$1 - \pi$

Table 2: The uncentred model

the plan to continue or not at an exit has no influence on the fact whether the driver is absentminded at a particular exit or not. Let $P(\omega)$ denote the prior probability that the driver assigns to state ω at the ex-ante stage. Then, from an ex-ante perspective it is optimal for the driver to leave the highway at an exit precisely when $4P(\omega_1) + P(\omega_{12}) \geq 1$, that is, when he deems it sufficiently likely that he will be absentminded at the first exit but not at the second one.

Now assume that $P(\omega_{12}) < 1$, the driver sees an exit, and reasons about whether to leave or not. Then he must conclude that he is not absentminded at both exits, and hence ω_{12} is no longer possible. If he uses Bayesian updating with respect to the event $E = \{\omega_0, \omega_1, \omega_2\}$ to revise his prior belief, then his updated belief $P(\cdot|E)$ will be given by

$$P(\omega_0|E) = \frac{P(\omega_0)}{1 - P(\omega_{12})}, P(\omega_1|E) = \frac{P(\omega_1)}{1 - P(\omega_{12})}, P(\omega_2|E) = \frac{P(\omega_2)}{1 - P(\omega_{12})}.$$

Hence, it will still be optimal to leave the highway at an exit precisely when $4P(\omega_1) + P(\omega_{12}) \geq 1$, just like at the ex-ante stage. In other words, there is no time inconsistency in this model.

The deeper reason for why the driver's preferences before and after seeing an exit are time consistent lies in the fact that contingent on being absentminded at both exits, his two acts would deliver exactly the same outcome: reaching the motel. Thus, by Savage's *sure-thing principle*, his preferences should not be affected by whether the state of being absentminded at both exits is ruled or not. This is in sharp contrast with many of the existing approaches, which *do* lead to decisions that are time inconsistent.

4. APPROACH WITH CENTRED POSSIBILITIES

The solution we discussed in the previous section relies on a reformulation of the absentminded driver problem. In particular, a difference is that in the original story, the driver knows that he will forget having passed any exit, even though he is aware of being at an exit at the time he reaches it. The second solution that we now turn to explore keeps the original formulation of the problem, and addresses one of the two concerns that we identified: namely, that in the formulation of the problem, the states are not mutually exclusive.

We begin by observing that there are three possible histories, terminating at the Motel (A), at Home (B), or at the Scary road (C), and the probability of each of these histories can be expressed in terms of the probability π that the driver continues at each time he arrives at an exit (Table 2).

motel		home		road	
<i>At1</i>	<i>At2</i>	<i>Bt1</i>	<i>Bt2</i>	<i>Ct1</i>	<i>Ct2</i>
$\alpha\pi^2$	$(1-\alpha)\pi^2$	$\beta(\pi-\pi^2)$	$(1-\beta)(\pi-\pi^2)$	$\gamma(1-\pi)$	$(1-\gamma)(1-\pi)$

Table 3: The centred worlds model

If the utility of each history is: $U(A) = 1$, $U(B) = 4$, $U(C) = 0$, then the expected utility is maximised at $\pi = \frac{2}{3}$. However, this way of representing the problem does not take into account the driver’s uncertainty about what the time period is (and therefore what exit he is approaching). But, of course, this makes a difference to what the driver should do: if he is at the first exit, then he should prefer to continue, while if he is at the second exit, then he should exit. We cannot express what exit the driver is approaching using the model in Table 2, since that does not model the different time periods at which the driver could be located. In other words, the simple model in Table 2 does not capture one important dimension of the driver’s uncertainty, that is which time period he is at.

To address this issue, we enrich the model to add the driver’s time location. This will allow us to represent the driver’s uncertainty regarding both the history and the time period at which he is located. The resulting refined model is summarised in Table 3.

As before, π represents the probability that the driver continues when reaching an exit, which we call the driver’s “choice probability”, and it determines the overall probability of each possible complete history. In addition to the choice probability, we add parameters to capture the driver’s uncertainty about the time period. α , β and γ stand for the probability that it is time 1, given that we are in history A , B or C , respectively. We can imagine that this is set on the basis of some randomising mechanism by nature. From now on, unless otherwise noted, we will assume that $\alpha = \beta = \gamma = \frac{1}{2}$.³

When the driver approaches an exit, the only thing he learns is that he hasn’t reached the Scary Road; in other words, he learns that it is not the second time period in history C , as otherwise he would have already abandoned the highway and reached the Scary Road. Out of the six possible centred locations, only five are compatible with his current evidence as he approaches the exit. He could be at the first time period, in which case each of the three histories is possible, as he could ultimately reach any of the final destinations depending on his next choices. Or he may already be at the second time period, and in this case he’d be on the road to either Home or the Motel, but certainly not the Scary Road, since he would have already passed that exit. We can represent the driver’s evidence set, as he approaches an exit, as the set $E = \{At1, At2, Bt1, Bt2, Ct1\}$.

The model in Table 2 did not allow us to formulate questions concerning which

³The uniform assignment can be motivated by appeals to symmetry. See Section 3 of Milano (2020) for a related discussion.

time period the driver is located at. But relative to the model in Table 3 we can now calculate the probability that the driver is at the first time period, and so approaching the first exit:

$$Pr(t1|E) = \frac{Pr(t1\&E)}{Pr(E)} = \frac{1/2}{\frac{1-\pi}{2} + \pi} = \frac{1}{1 + \pi}$$

Note that $Pr(t1|E) = 1/2$ only if $\pi = 1$, so only if the driver is certain to continue at each exit the probability of being at the first one is $1/2$. But we cannot fix the probability of being at the first exit without knowing the value of π , given that the driver is absentminded.

As π represents the choice probability, which is decided when the driver reaches an exit, its value is fixed then. At the time of choice, that is when he reaches an exit and his evidence is E , the driver should fix a value of $\pi = x$ that maximises expected utility, given the payoffs associated with each history:

$$EU_E(\pi = x) = Pr_E(A|\pi = x)U(A) + Pr_E(B|\pi = x)U(B) + Pr_E(C|\pi = x)U(C)$$

To find the value of x that is optimal, we need to know the conditional probability of each history, conditional on fixing the value of π . We know the following:

$$\begin{aligned} Pr(At1|E) &= Pr(At2|E) = \frac{\pi^2}{\pi + 1}; \\ Pr(Bt1|E) &= Pr(Bt2|E) = \frac{\pi - \pi^2}{\pi + 1}; \\ Pr(Ct1|E) &= \frac{1 - \pi}{\pi + 1}; \text{ and } Pr(Ct2|E) = 0. \end{aligned}$$

From this, it follows that the expected utility above is maximised at $\pi = \sqrt{\frac{7}{3}} - 1$, or approximately 0.53. This value differs from most other solutions in the literature, and is also different from the value of π according to the *ex ante* optimal strategy, which remains $\frac{2}{3}$ also in our approach.

Before moving on, two remarks will be worth making. We could in fact recover the $\frac{2}{3}$ solution in our framework in at least two ways, that however have independent drawbacks (on which more below). First, we could recover the $\frac{2}{3}$ solution if, instead of Bayesian updating, we calculated the driver's posterior probabilities when reaching an exit by using imaging (Lewis, 1976; Cozic, 2011). In this case, upon seeing an exit, the prior probability of $Ct2$ is shifted to the "most similar" state that is still possible, which is $Ct1$. As a result, the driver's expected utility is maximised at $\pi = \frac{2}{3}$. In fact, this is independent of how the parameters α , β and γ are chosen. A second way in which the $\frac{2}{3}$ solution could be recovered in our model is by changing the value of the parameter γ to 1. This reasoning would be analogous to Lewisian halving in the Sleeping Beauty problem (Milano, 2020), and we will not rehearse here arguments supporting it. However, a significant problem with both these approaches would be to open the door to diachronic inconsistency, as we argue in the next section.

	safe harbour	scary road
bet B_1	a	b
bet B_2	c	d
total payoff	$a + 2c$	$b + d$

Table 4: A Dutch book for the absentminded driver problem

5. DUTCH STRATEGIES

In this section we show that, within the centred worlds model of Section 4, Bayesian updating is both necessary and sufficient for being invulnerable against diachronic Dutch strategies.

5.1. DEFINITION OF DUTCH STRATEGIES

In general, a *Dutch strategy* is a system of bets that guarantees a sure loss from an ex-ante perspective, but that the decision maker is nevertheless willing to accept, given his conditional beliefs, upon receiving new information. It thus points at a “discrepancy” between the decision maker’s ex-ante beliefs and his conditional beliefs. For traditional decision problems, Dutch strategies are typically called *Dutch books*, and for those scenarios it is well-known that a Dutch strategy will never be accepted precisely when the decision maker revises his beliefs by Bayesian updating (Teller, 1973; Skyrms, 1987).

The absentminded driver problem, as we have seen, is not a traditional decision problem. However, Dutch strategies can still be defined for this type of problem (Hitchcock, 2004). For example, consider a system of two bets, where both bets concern the event of reaching a safe harbour (the motel or the driver’s home) versus reaching the scary road. The first bet, B_1 , is offered ex-ante, whereas the second bet, B_2 , is offered after the driver leaves the friend’s house, every time he reaches an exit. Note that, if a safe harbour is eventually reached, then the bet B_2 will be offered twice on the highway without the driver knowing it. Indeed, if the driver would realise that the bet B_2 is offered for the second time on the highway, then he would know that he is at the second exit, contradicting the assumption that he is absentminded. As such, when deciding whether or not to accept the bet B_2 , the driver focuses only on the expected payoff from accepting this bet now, since he has no immediate control over his betting decision at the other time period. Such a system of bets can be summarised by Table 4.

Here, a, b, c and d are (possibly negative) payoffs. Hence, if the driver accepts the bet B_1 , then he receives a if a safe harbour is reached, and b if the scary road is reached. Similarly, if he accepts the bet B_2 , then he receives c if a safe harbour is reached, and d if he reaches the scary road. If the payoff is negative, say $-e$ with $e > 0$, then by receiving this payoff we actually mean that the driver must pay the amount e . Recall that, if a safe harbour is reached, then the bet B_2 will be offered twice, at both exits of the highway. Hence, if you accept B_2 at both occasions, the amount c will be received twice too, resulting

in a total amount of $2c$. This system of bets is called a *Dutch strategy* if $a + 2c < 0$ and $b + d < 0$. In this case, the total payoff from accepting the two bets, no matter whether you reach a safe harbour or the scary road, will always be negative.

When will the driver be willing to accept the bets B_1 and B_2 ? Precisely when the expected payoff from accepting the bet is greater than zero. It is important that the bet B_1 is evaluated with respect to the ex-ante beliefs, whereas the driver uses his conditional beliefs to evaluate the bet B_2 . Suppose that the driver, ex-ante, assigns probabilities p and $1 - p$ to reaching a safe harbour and reaching the scary road, respectively. Then, his expected payoff from accepting the bet B_1 would be

$$EU(B_1) = p \cdot a + (1 - p) \cdot b.$$

Assume that his conditional beliefs, upon seeing an exit on the highway, assign probabilities q and $1 - q$ to these two events. Then, the expected payoff from accepting the bet B_2 on the highway would be

$$EU(B_2) = q \cdot c + (1 - q) \cdot d.$$

Therefore, the driver would be willing to *accept* the system of bets if $EU(B_1) > 0$ and $EU(B_2) > 0$.

5.2. IMMUNITY AGAINST DUTCH STRATEGIES

We will now show that the driver is immune against Dutch strategies precisely when he revises his beliefs by Bayesian updating. Before doing so, we first characterise Bayesian updating in terms of the relationship between the ex-ante probability p and the conditional probability q assigned to a safe harbour.

As before, let the three destinations *motel*, *house* and *scary road* be denoted by A , B and C , respectively, and let $At1$, $At2$, $Bt1$, $Bt2$, $Ct1$, $Ct2$ be the centred possibilities of reaching each of these destinations, indexed by time. We have seen that seeing an exit can be identified with the event

$E = \{At1, At2, Bt1, Bt2, Ct1\}$, where the possibility $Ct2$ has been ruled out. If we adopt the notation from Table 3, with $\alpha = \beta = \gamma = \frac{1}{2}$, then the ex-ante probability assigned to a safe harbour is

$$p = \pi$$

whereas the conditional probability assigned to a safe harbour, obtained from Bayesian updating after seeing an exit, is

$$q = \frac{\pi}{1 - \frac{1}{2}(1 - \pi)} = \frac{2p}{1 + p}.$$

Therefore,

$$\frac{q}{1 - q} = 2 \cdot \frac{p}{1 - p}.$$

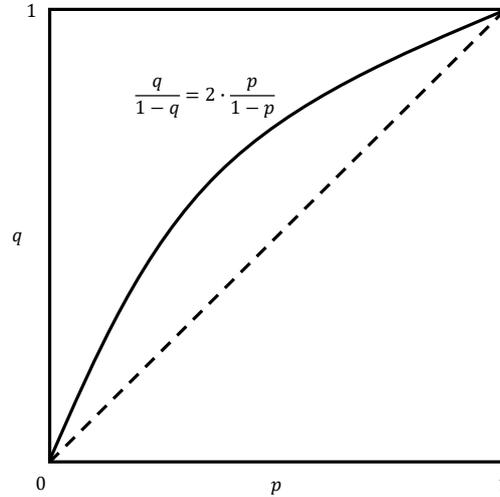


Figure 5.1: Immunity against Dutch strategies

This relationship characterises Bayesian updating.

The following result states that the driver will only be immune against Dutch strategies if he revises his beliefs by Bayesian updating.

Theorem 5.1. Suppose that ex-ante, the driver assigns probability p and $1 - p$ to reaching a safe harbour versus reaching the scary road, and that on the highway he assigns probabilities q and $1 - q$ to these two events. Assume that both p and q are strictly between 0 and 1. Then, there is no Dutch strategy of the form above that the driver is willing to accept, if and only if,

$$\frac{q}{1-q} = 2 \cdot \frac{p}{1-p},$$

that is, if and only if the driver revises his beliefs by Bayesian updating.

The condition in this theorem is visualised in Figure 5.1, and the formal proof is given in the Appendix. The theorem thus asserts that, if we adopt the model with centred possibilities as proposed in Section 4, then the solution we put forward in that section is the only plan that (a) is optimal for the driver once he sees an exit, and (b) makes him immune against Dutch strategies. Any other belief revision rule, including imaging, allows for the design of a Dutch strategy that the driver is willing to accept.

6. TIME CONSISTENCY

A common concern in the analysis of the absentminded driver problem is time consistency. As we recalled in the introduction, a recurring motivation for the solutions that have been discussed in the literature is to resolve the apparent inconsistency between the driver's optimal strategy *ex ante* (before he starts the journey) and when he reaches an exit. This poses a problem because it is implicitly assumed that the driver does not learn anything new between these times – call this the *no news assumption*. He always knew that he was going to reach an exit, and he should not be surprised when he does. Therefore, the driver should not update his beliefs upon reaching an exit, and the discrepancy between his optimal strategies when considered *ex ante* vs. when reaching an exit (what we will call *interim*) cannot be explained by a change in the probabilities assigned to the relevant states of the world. Starting from the *no news assumption*, most of the solutions in the literature (see Section 2) attempt to resolve the issue of time inconsistency by looking for a strategy that could be characterised as rational, according to some modified notion of optimality, both *ex ante* and *interim*.

Contrary to this received approach, our analysis highlights that the *no news assumption* is incorrect. The driver learns some new relevant information upon arriving at an exit, namely that *Ct2* is not the case at this point. Since this is a possibility that the driver eliminates when he reaches an exit, he should update on this information. Failing to do so, moreover, would lead to diachronic inconsistency and vulnerability to a Dutch strategy. Therefore, the driver's beliefs *ex ante* and upon reaching an exit differ, but are not time inconsistent. On the contrary, the difference results from applying standard Bayesian updating upon learning a new piece of evidence. Our analysis explains why the driver's changing beliefs, and the resulting change in optimal strategies, are in fact time consistent, eliminating the need to resolve the apparent discrepancy that motivated previous solutions presented in the literature.

Still, one may pose the following objection to our analysis, based on the idea of forming a commitment. Suppose that the driver has the possibility, *ex ante*, of forming a commitment, that is, selecting a strategy that will be binding for him whenever he reaches an exit. Based on our analysis, the optimal strategy for the driver *ex ante* would be to continue with a probability of $\frac{2}{3}$. However, if he were to commit to this strategy, he would regret doing so at the moment he reaches an exit, when the optimal strategy after updating would be to continue with probability 0.53. Under these conditions, would it be rational for the driver to choose to commit to the *ex ante* optimal strategy? Or should he commit to the strategy he knows will be optimal for him after updating, even though it is not optimal at the time he makes the commitment?

Neither option seems rational. If the driver commits to the *ex ante* optimal strategy, he will come to regret it, based on his own beliefs at the time when he needs to implement the strategy. Choosing to commit to the other strategy, on the other hand, would be irrational from the *ex ante* perspective. Thus,

if the driver has the option to commit, it seems that neither of the available strategies (the *ex ante* optimal one or the *interim* optimal one) are rational to commit to.

For a rational agent who plans to update her beliefs via conditioning, and does not expect to forget any of her experiences, forming a commitment is redundant, since her present and future (expected) optimal strategies are identical (see, for instance, Perea, 2012, Lemma 8.14.9). However, if the agent expects her future self to become irrational, then it may be rational to form a commitment now in order to prevent herself from making a bad choice in the future, as in the story of Ulysses tying himself to the mast of his ship in order to escape the sirens' call (van Fraassen, 1995). Thus, forming a commitment in the face of an expected deviation, in the future, from what is the present optimal strategy, may be rational if the agent expects to become irrational by the time of making a decision.

Interestingly, the case of absentmindedness illustrates how forming a commitment is not always either redundant or necessary to prevent anticipated irrationality. Sometimes, it is simply not possible to rationally commit. Although his future beliefs deviate predictably from his *ex ante* beliefs, the driver's future self is not irrational. Upon reaching an exit, the driver arrives at the new beliefs by conditioning on newly acquired evidence, as we have seen in Section 4, and committing to the *ex ante* optimal strategy would make him vulnerable to a Dutch Strategy, as we have shown in Section 5. So, although the driver expects that he will deviate from his *ex ante* optimal strategy when he will make a decision in the future, he can't justify committing to the *ex ante* optimal strategy on the basis of his own future irrationality.

Absentmindedness means that the future driver cannot locate himself in time. When he approaches an exit, he is uncertain whether it is the first or the second one. This happens predictably and therefore his evidence at any time that he approaches an exit is the same. But this also means that when he considers, *ex ante*, what his own beliefs will be *later* during the journey, on the highway, he is not referring to his future self at a definite point in time. *Later* here refers both to his future self approaching the first exit, and his future self approaching the second exit (if he gets that far).⁴ If the driver were expected to defer to his rational, future self, this would raise the question: to his future self *at which point in time?*

Ex ante, the driver can only be certain that he will hit the first exit (since that exit is reached no matter what the final destination), at which point he will update his beliefs and optimal strategy in the way that we have described. So, a natural answer to the previous question is to defer to the driver's future self at the only time he is certain to reach an exit, that is t_1 . Interpreted in this way, *later* picks out t_1 , but it is not a *stopping time* (Schervish et al., 2004).

⁴Since it does not pick out a unique time for each of the possible histories, *later* in this loose sense is not a well-defined random time as discussed by Schervish et al. (2004).

A (possibly random) time T is called a *stopping time* when, at any point in time, the agent knows whether T has already occurred or not. An example of a stopping time would be the day of your son's first birthday (assuming that you do not forget what day it is!), or the first time a coin that you are about to toss repeatedly comes up Heads, assuming that you observe and don't forget any of the previous tosses. Since the driver does not know he has reached the first exit when he is there, t_1 is not a stopping time in our case. Moreover, learning that *later*, so interpreted, has arrived would give the driver information that changes his optimal strategy: if he knew that he was approaching the first exit, then it would be optimal for him to continue.

Under these conditions (that is, *later* not being a stopping time, and the additional information that *later* has arrived being relevant to the driver's other beliefs), Schervish et al. (2004) show that the so-called Principle of Reflection does not apply.⁵ This principle, introduced by van Fraassen (1984), says that for any two times $t_1 < t_2$, and for any event A , your credence in A at t_1 , conditional on your credence in A at t_2 being r , should also be r . That is: $P_{t_1}(A|P_{t_2}(A) = r) = r$. Since Reflection does not apply to our case, the absent-minded driver's *ex ante* beliefs do not reflect his expected *later* beliefs (Schervish et al., 2004). Therefore, he cannot rationally commit *ex ante* to a strategy that is optimal *later*, since his own beliefs *ex ante* should not rationally reflect the beliefs he will have *later*.

Anticipating his own absentmindedness, the driver can only make rational decisions at the moment when it is required. If he is required to pick a strategy *ex ante*, and prevented from changing it along the way, then he should pick the *ex ante* optimal strategy, which differs from what he would choose on the highway. In this scenario, the driver is effectively occupying the role of a passenger, since he has no control over the direction travelled by the car during the journey. If, instead, the driver is allowed to make choices while on the highway, then he should opt for the *interim* optimal strategy, deviating from the choice probability that he would have picked *ex ante*. This is not a sign of dynamic inconsistency, since the change in beliefs is the result of rational updating, and commitment is not appropriate in this case.

7. CONCLUSION

The *absentminded driver problem* presents a decision situation that challenges standard approaches to modelling one-person decision problems. First, act-state independence, a core assumption of the Savage (1972) framework, is violated. Second, the decision problem appears unstable, since the optimal choice appears to change if we consider it from an *ex ante* perspective, or at the moment of choice, resulting in time inconsistency.

We have presented two new approaches to the resolution of the absentminded driver problem that address these two issues. In the first proposal, we present

⁵For further discussion of the limitations of the Reflection principle, see also Mahtani (2017).

a version of the absentminded driver problem where we model the driver as being uncertain about his own degree of absentmindedness. This approach is completely in line with Savage’s framework, and the resulting solution satisfies both act-state independence and time consistency.

In the second proposal, we model states as centred possibilities. This allows us to model explicitly the driver’s uncertainty regarding both the destination that he will eventually reach, and the time period at which he is located. Using this model, we derive the optimal choice probability of continuing for the driver, taking into account the centred evidence that is available to him at the time he makes a choice. We show that Bayesian conditioning in this setting is the only updating strategy that prevents vulnerability to a sure loss via a Dutch strategy.

A surprising implication of our second approach is that it highlights how commitment is sometimes impossible, since the driver’s *ex ante* and *interim* optimal strategies differ predictably but rationally. We have argued that this conclusion does not violate time consistency, because it is based on updating on new evidence. Our discussion of time consistency in the setting of absentmindedness has also highlighted the relationship between planning, commitment and van Fraassen’s Reflection principle. *Ex ante* planning is not possible, because commitment to any strategy but the one that is optimal at the time of making a choice is not appropriate in similar cases, for wholly anticipated and yet rational reasons.

REFERENCES

- Aumann, R. J., Hart, S., and Perry, M. (1997). The Absent-Minded Driver. *Games and Economic Behavior*, 20(1):102–116.
- Cozic, M. (2011). Imaging and Sleeping Beauty: A case for double-halfers. *International Journal of Approximate Reasoning*, 52(2):137–143.
- Gilboa, I. (1997). A Comment on the Absent-Minded Driver Paradox. *Games and Economic Behavior*, 20(1):25–30.
- Halpern, J. Y. and Pass, R. (2016). Sequential equilibrium in games of imperfect recall. In *Proceedings of the Fifteenth International Conference on Principles of Knowledge Representation and Reasoning*, KR’16, pages 278–287, Cape Town, South Africa. AAAI Press.
- Hitchcock, C. (2004). Beauty and the Bets. *Synthese*, 139(3):405–420.
- Jeffrey, R. C. (1990). *The Logic of Decision*. University of Chicago Press, Chichago, 2nd edition edition.
- Lewis, D. (1976). Probabilities of Conditionals and Conditional Probabilities. *The Philosophical Review*, 85(3):297–315. Publisher: [Duke University Press, Philosophical Review].
- Lewis, D. (1981). Causal decision theory. *Australasian Journal of Philosophy*, 59(1):5–30. Publisher: Routledge _eprint: <https://doi.org/10.1080/00048408112340011>.
- Mahtani, A. (2017). Deference, respect and intensionality. *Philosophical Studies*, 174(1):163–183.

- Milano, S. (2020). Bayesian Beauty. *Erkenntnis*.
- Perea, A. (2012). *Epistemic Game Theory: Reasoning and Choice*. Cambridge University Press.
- Piccione, M. and Rubinstein, A. (1997). On the Interpretation of Decision Problems with Imperfect Recall. *Games and Economic Behavior*, 20(1):3–24.
- Rabinowicz, W. (2003). Remarks on the Absentminded Driver. *Studia Logica*, 73(2):241–256.
- Savage, L. J. (1972). *The Foundations of Statistics*. Courier Corporation. Google-Books-ID: zSv6dBWneMEC.
- Schervish, M. J., Seidenfeld, T., and Kadane, J. B. (2004). Stopping to Reflect. *The Journal of Philosophy*, 101(6):315–322.
- Schwarz, W. (2015). Lost memories and useless coins: revisiting the absent-minded driver. *Synthese*, 192(9):3011–3036.
- Skyrms, B. (1987). Dynamic Coherence and Probability Kinematics. *Philosophy of Science*, 54(1):1–20. Publisher: The University of Chicago Press.
- Skyrms, B. (1990). *The Dynamics of Rational Deliberation*. Harvard University Press.
- Teller, P. (1973). Conditionalization and observation. *Synthese*, 26(2):218–258.
- van Fraassen, B. C. (1995). Belief and the problem of Ulysses and the sirens. *Philosophical Studies*, 77(1):7–37.
- van Fraassen, C. (1984). Belief and the Will. *The Journal of Philosophy*, 81(5):235–256. Publisher: Journal of Philosophy, Inc.

A. APPENDIX

Proof of Theorem 5.1.

We first show that, if $\frac{q}{1-q} = 2 \cdot \frac{p}{1-p}$, then the driver will never accept a Dutch strategy. To see this, suppose that the driver would be willing to accept the bets B_1 and B_2 . Then, $EU(B_2) > 0$ and hence

$$q \cdot c + (1 - q) \cdot d > 0.$$

Dividing by $1 - q$ thus yields

$$\frac{q}{1 - q} \cdot c + d > 0.$$

As $\frac{q}{1-q} = 2 \cdot \frac{p}{1-p}$, we obtain

$$\frac{p}{1 - p} \cdot 2c + d > 0,$$

and hence

$$p \cdot 2c + (1 - p) \cdot d > 0.$$

Moreover, as the driver is willing to accept the bet B_1 , it must be that $EU(B_1) > 0$, which means

$$p \cdot a + (1 - p) \cdot b > 0.$$

By adding up the last two inequalities we get

$$p \cdot (a + 2c) + (1 - p) \cdot (b + d) > 0,$$

which implies that either $a + 2c > 0$ or $b + d > 0$. Thus, the system of bets that the driver is willing to accept cannot be a Dutch strategy.

Next, assume that the condition in the theorem is violated, that is, $\frac{q}{1-q} \neq 2 \cdot \frac{p}{1-p}$. We will distinguish two cases: (1) $\frac{q}{1-q} < 2 \cdot \frac{p}{1-p}$, and (2) $\frac{q}{1-q} > 2 \cdot \frac{p}{1-p}$. For both cases we will construct a Dutch strategy that the driver is willing to accept.

Case 1. Assume that $\frac{q}{1-q} < 2 \cdot \frac{p}{1-p}$.

Define the number

$$\varepsilon := \frac{2p/(1-p)}{q/(1-q)} - 1,$$

which is positive. Let the payoffs of the bets B_1 and B_2 in Table 4 be given by

$$a := 2, \quad b := -(1 + \frac{3}{4}\varepsilon) \cdot \frac{q}{1-q}, \quad c := -1 - \frac{1}{4}\varepsilon, \quad d := (1 + \frac{1}{2}\varepsilon) \cdot \frac{q}{1-q}.$$

We show that this system of bets is a Dutch book that the driver is willing to accept. First, it may be verified that $a + 2c < 0$ and $b + d < 0$, and therefore this is a Dutch strategy. Moreover,

$$\begin{aligned} EU(B_1) &= p \cdot 2 - (1-p) \cdot (1 + \frac{3}{4}\varepsilon) \cdot \frac{q}{1-q} \\ &> p \cdot 2 - (1-p) \cdot (1 + \varepsilon) \cdot \frac{q}{1-q} \\ &= p \cdot 2 - (1-p) \cdot 2 \cdot \frac{p}{1-p} \\ &= p \cdot 2 - p \cdot 2 = 0. \end{aligned}$$

In the second equality we have used the definition of ε . Thus, the driver is willing to accept B_1 . Also,

$$\begin{aligned} EU(B_2) &= -q \cdot (1 + \frac{1}{4}\varepsilon) + (1-q) \cdot (1 + \frac{1}{2}\varepsilon) \cdot \frac{q}{1-q} \\ &= -q \cdot (1 + \frac{1}{4}\varepsilon) + q \cdot (1 + \frac{1}{2}\varepsilon) > 0, \end{aligned}$$

which means that the driver is also willing to accept B_2 . Summarising, we see that there is a Dutch strategy that the driver is willing to accept.

Case 2. Suppose that $\frac{q}{1-q} > 2 \cdot \frac{p}{1-p}$.

Define the number

$$\varepsilon := 1 - \frac{2p/(1-p)}{q/(1-q)},$$

which is positive. Set the payoffs of the bets B_1 and B_2 in Table 4 by

$$a := -2, \quad b := \left(1 - \frac{3}{4}\varepsilon\right) \cdot \frac{q}{1-q}, \quad c := 1 - \frac{1}{4}\varepsilon, \quad d := \left(-1 + \frac{1}{2}\varepsilon\right) \cdot \frac{q}{1-q}.$$

Like in Case 1, we show that this system of bets is a Dutch strategy that the driver is willing to accept. First, it holds that $a + 2c < 0$ and $b + d < 0$, and therefore this is a Dutch strategy. Also,

$$\begin{aligned} EU(B_1) &= -p \cdot 2 + (1-p) \cdot \left(1 - \frac{3}{4}\varepsilon\right) \cdot \frac{q}{1-q} \\ &> -p \cdot 2 + (1-p) \cdot (1-\varepsilon) \cdot \frac{q}{1-q} \\ &= -p \cdot 2 + (1-p) \cdot 2 \cdot \frac{p}{1-p} \\ &= -p \cdot 2 + p \cdot 2 = 0. \end{aligned}$$

In the second equality we have used the definition of ε . Hence, the driver is willing to accept B_1 . Moreover,

$$\begin{aligned} EU(B_2) &= q \cdot \left(1 - \frac{1}{4}\varepsilon\right) + (1-q) \cdot \left(-1 + \frac{1}{2}\varepsilon\right) \cdot \frac{q}{1-q} \\ &= q \cdot \left(1 - \frac{1}{4}\varepsilon\right) + q \cdot \left(-1 + \frac{1}{2}\varepsilon\right) > 0, \end{aligned}$$

which means that the driver is also willing to accept B_2 . Hence also in Case 2 we see that there is a Dutch strategy that the driver is willing to accept.

By combining Cases 1 and 2 we conclude that, whenever $\frac{q}{1-q} \neq 2 \cdot \frac{p}{1-p}$, we can always construct a Dutch strategy that the driver is willing to accept.