A PROBLEM CONCERNING CHANGING PREFERENCES – 30 June 2003

My sole aim in this paper is to point to a problem in the area of intertemporal choice. The problem concerns the justifiability of having taken a certain course of action, once it was completed. Throughout I will suppose that this course was adopted on the basis of principles of rational choice, as available to the agent when he was still deliberating about what to do. My interest is in the claim that given this supposition, the decision for the course of action is also always justifiable ex post.

There are many ways in which the decision-maker might become disappointed with an earlier decision he made. For example, bad chance may lead to an unsatisfactory outcome even if he followed a strategy that he regarded as optimal at the outset. There is then room for at least some regret. However, if while assessing the reasons for this bad outcome, this agent could also be prompted to revisit the considerations that made him adopt his strategy and might well ascertain that he could not have found a better plan than what he actually pursued. This will be the case if he can assure himself that he was making his decision in light of principles of rational choice. Then he has this thought available: if he was to make the decision again, he would do the same thing. In this sense, the decision is justified.

I will not only present decision-makers who follow principles of rational choice, but I will also only consider choice problems that can be described and analyzed by standard decision theory and its straightforward extensions. Furthermore, I will try to follow the usual prescription this theory issues for solving these problems. So my question, more precisely, is this: is it true that if agents follow the prescriptions of rational choice theory in intertemporal decision problems then no matter how things transpire they will always feel justified once the problem has unfolded? For an answer, I will develop a problem case that indicates that this claim is not to be accepted.

Section I below describes the canonical model for intertemporal choice and explains how the issue of ex post justification arises in it. Section II treats decision problems with changing preferences. Then section III presents the problem case in which, as I will argue, ex post justifiability is not always available. In section IV, I will discuss another way in which changing reasons in decision problems appear, that of losing information within them.

I

Consider the following decision problem (Example 1). At t_1, an agent can choose between two actions, x and y. If he chooses y, he will get payoffs at the level of 2 at t_3. If he chooses x, he will have to make a further decision at t_2, between z and w. If he there chooses z, the eventual payoffs at t_3 will be 4, alternatively (by choosing w), he will get 0. If we can suppose that the preferences of the agent that are represented by the payoff numbers just given are rational, it seems fairly clear that this decision-maker should settle at t_1 to choose x first and z after that. Moreover, neither at t_2 nor at t_3 ought he to think that this plan is up to criticism or ought to have been altered.

One of the reasons why this analysis seems so trivial is the sheer simplicity of the decision problem. Not only does it unfold in two periods, but it has no chance move in it, nor does it exhibit changing preferences or imperfect recall. As I will shortly present and discuss problems for which one or more of these characteristics do not hold, it is worthwhile to describe the generic model of intertemporal decision situations to which each of these will conform.

I will suppose that a finite decision problem is spanned by a decision tree, will refer to the nodes of this tree as the set of histories in the problem.[1] Let H denote this set, the set of sequences of events (actions or chance moves) in the decision problem starting from the root of the tree. Histories that are not followed by further events are terminal histories, and belong to the set Z. The set C stands for the histories where chance makes a move. Then the set of histories D=H-Z-C are those where the decision-maker is to act. A decision problem also provides information about the probabilities with which chance makes this or that move at histories in C. The function $u_h: Z \rightarrow \Re$ specifies the payoff numbers the agent attaches to each terminal history at a history h. Finally, E denotes the set of information sets in the problem, being a partition of the set D. Decision problems might be sorted out into the following classes. In problems with perfect information each information set is singleton. In problems without chance moves C is the empty set. In problems without preference change $u_h=u$ at each history h in D. And let us only say here that in problems with perfect recall the partition E has a given and well-specified structure. In this paper I will only discuss problems with perfect information and prefect recall, so information sets will play no role. Let me call problems that in addition feature no changing preferences straightforward problems.

Solutions of decision problems are expressed as (pure) strategies that assign to each history in D a specific action. The set of actions that actually have to be made according to a strategy might be called plans or courses of action. The agents are motivated by the aim to maximizes their expected payoffs. In straightforward problems this means that they are to find the best plan, that is the plan that maximizes their ex ante expected payoffs. Such a plan is generated by an optimal strategy. Note that the analysis of these problems is much simplified by considering the following claim. Define behavioral strategies as those that render a probability distribution over each action following a history, for each history in C. Then it is true for straightforward problems that no behavioral strategy can improve on an optimal strategy.

Let us turn to the issue of ex post justifiability of a course of actions, decided at the outset of the decision problem. The notions of a strategy does not contain a view about whether indeed it will be followed in a given problem. Suppose for now that the agent can always assess at each point he has to make a move whether it is worthwhile to act as the strategy prescribes him to act there. Then we may say that at these points the strategy is justified for him if he cannot find a better course of actions from that point on to follow. This kind of ex post justifiability is often referred to as time consistency, and may be formalized as follows. Let s* be an optimal strategy, and s any other strategy, then in straightforward problems, and for each history that is reached with positive probability given s*, s* is time-consistent if:

(1) $\sum_{z\in\Sigma} p(z|h,s^*)u(z) \geq \sum_{z\in\Sigma} p(z|h,s)u(z)$

(Here $p(z|h,s)$ is the probability of getting to z from h given strategy s.) We can also claim that in these problems an optimal strategy is always time-consistent, as optimal strategies s* here can be defined as (∅ denotes the root history):

(2) $\sum_{z\in\Sigma} p(z|\varnothing,s^*)u(z) \geq \sum_{z\in\Sigma} p(z|\varnothing,s)u(z)$

---

[1] The description of intertemporal decision problems owes much to chapter 4 of Ariel Rubinstein's Modeling Bounded Rationality (MIT Press, 1998).

Now how can the agent evaluate the strategy he has followed when the decision problem is completed? Here we cannot employ the notion of time consistency as there is no further move to make. Indeed, time consistency is a forward-looking notion. But one may seek to justify a course of action that has unfolded, as well. Such justification may refer to its time consistency at the relevant histories, but as we have seen this thought might well lead in straightforward problems to ascertaining that the course of action was ex ante optimal.

Now is it true that even if luck was against the agent, he could ascertain that his deliberation and subsequent choice was justified? Suppose then that an agent faces the following decision problem (Example 2). At the outset (at t_1), he could again choose between two courses of action, x and y. If he chooses the latter option, he will gain payoffs of 2 at t_3. If he opts for x, what he will get at t_3 depends on the realization of a possibility at t_2. Suppose also that the agent does not have to undertake any further action between t_2 and the time he receives his gains. If the possibility will be indeed realized, that is if θ holds, then he will gain 4; otherwise he will gain 0. Denote the probability that the possibility will be realized with p. Then setting aside the issue of risk aversion, we will have to say that the agent should choose x if and only if p is larger than one-half.

The possibility that had not unambiguously declared itself at t_1 is regarded by decision theory as an event outside of the control of the agent. It could for instance be related to a prevailing level of market prices, or just to the outcome of a lottery draw, or to weather conditions. Consider now the scenario when p is larger than one-half, then at t_1 x should be taken. Suppose next that the possibility was not realized ($\sim\theta$ holds). Should the agent regret this decision? As the outcome is a gain of 0, whereas having opted for y he could have got 2, he certainly should feel some regret.

Given the parsimonious structure of Example 2, time consistency cannot really be tested in it. Of course, this example can be easily extended so that the agent is to take actions after the value of θ is learned (but before t_3). But one can reproduce the question about ex post justifiability at terminal nodes in such kind of an extension as well. Consider then that this example exhibits a straightforward decision problem as these were defined above. In these problems optimal strategies are also time-consistent. Then we could say that in these problems, at terminal nodes actually reached, the agent finds the strategy he followed justified, if it was optimal. As the strategy in Example 2 was optimal, it is to be regarded as justified even if it yielded

0. Let us say that this is because at the time of the decision all the reasons the agent could have had were on its side.

II

The two examples and their analyses in the previous section were very familiar. Decision problems in which the agent can foresee that his preferences will change at a later time pose more difficulties (Example 3-A). At the outset (at t_1), again, he could choose between two courses of action, x and y. If he chooses the latter option, he will get payoffs of 2 at t_3, no matter what. If he opts for x, he can make an other decision at t_2, between z and w. According to his preferences at t_1, if he takes z he will gain 4; otherwise he will gain 0. But he can also foresee that his preferences will be different at t_2: he will think then that opting for z yields 0, whereas opting for w gives him 1. If the agent would like to act according to his preferences at t_1, then he clearly should choose x, followed by z. Can this strategy be conceived as justifiable ex post?

Ex post justification was suggested to be related to time consistency. Now the strategy (x,z) is not time-consistent in the sense that the agent at t_2 will wish to do something else than z (namely w). As preferences change here, the definition of time consistency as defined by inequality (1) cannot be used here. Instead we may stipulate that a strategy s* is time-consistent if for each history h that is reached with positive probability according to it, and for any other strategy s

(3) $\sum_{z\in\Sigma} p(z|h,s^*)u_h(z) \geq \sum_{z\in\Sigma} p(z|h,s)u_h(z)$

In Example 3-A, again, this inequality does not hold for the history generated by action x.

As below, we may think of a course of actions being justifiable ex post, if it is also ex ante optimal. But it is not clear what ex ante optimality means here. The course (x,z) can be regarded as ex ante optimal, given preferences at t_1. But it might not be carried out. Consider however a version of the above problem in which preferences at t_2 change so that if z is chosen 1 will be had, and w gives the agent 0 (Example 3-B). In this case, the expected t_2 choice is z, and then from the point of view of t_1 the highest possible payoffs will be gathered. Now is this strategy justifiable from the point of view of t_3? Given the preferences the agent has at this point, he will feel that the option for y at t_1 would have been better. It is certainly not true that all the reasons were on its side.

But how can the agent deal with his conflicting reasons? In order to be able to make then an assessment of which strategy could be regarded as justifiable ex post, we may decide to emphasize that the agent cannot have an influence on what he will do at a later time solely by means of determining a strategy, or by just wishing that his later self conforms to his present intentions. Agents always act according to their actual preferences, so we can think of intertemporal decision situations as a game played by the various selves of the agent identified by histories in the decision problem. In straightforward problems without preference change, ex ante optimal strategies will induce a subgame-perfect Nash-equilibrium of the game between selves. But in decision problems with preference change, there is no guarantee that this equivalence holds. Indeed, the subgame-perfect Nash-equilibrium in Example 3_A calls for playing y at t_1, yielding a payoff of 2; whereas the optimal strategy based on the t_1 preferences of the agent promises 4. Alternatively, ex ante optimality could be defined as whatever the game between the selves brings about. As Example 3-B shows, this might well not be justifiable for the agent at t_3, again, as we cannot presume that all the reasons are on its side.

The agent in these examples has two sets of preferences issuing different reasons for actions. But why do we care that much about the ones appearing later in time from the point of view of ex post justifiability? After all, it could be the case that the preferences in Example 3-A, say, are due to the agent being tired or simply distracted by some features not specified in the decision situation. Similarly, it could be the case that, in the same example, the agent gives in to some sudden outburst of enthusiasm at t_1, or is just not careful enough to consider how he will assess the problem in the middle of executing it. Or he is afraid of a new experience, and so on.

In decision theory it is often acceptable to regard the preferences one has at any history of the decision situation as providing sufficient grounds for acting in this or that way. In straightforward problems these assumptions might indeed be unproblematic to make, as it is that preferences provide no reasons while choosing a strategy. But in problems that feature changing preferences, and where preferences at each history are determined, it is an issue which preferences are justified to have. Some such preferences may not be relevant or authoritative, in the sense that reason does not support them.

Indeed, I would like to propose that rational deliberation in an intertemporal decision problem calls for a decision about which preferences are to be acted upon. Rationality of preferences does not have to mean only that at any given history of the problem the agent’s preferences are coherent. Note that this proposal calls for giving up a basic tenet of decision theory, namely, that whatever is to motivate the agents to act is given by the description of the problem itself as a model. This issue is independent from that of the capacity to control one’s behavior within the decision problem. In my view whether the agent has such a capacity or not is not a subject for rational choice theory; not acting upon a rational strategy is a hallmark of irrationality. Indeed I will assume that agents have a second-order desire to carry out strategies reached by rational deliberation, in intertemporal decision problems.

But how could agents form a rational course of action for problems which exhibit changing preferences? Without pretending to provide here a full answer to this question, let me outline a procedure the following of which could easily yield a determinate overall preference ordering in such problems. As the decision about a strategy is to made at the outset of a decision problem, the agent will be required to make judgements about the status of his future preferences by fully imagining, as much as this is possible, the authority of the various preferences that might arise in the given decision problem and what it would be like to act upon them. He may go through the following steps:

(a) In some cases, certain preferences can be eliminated on the ground of their being formed under unfavorable circumstances, or being adopted because of some palpably false beliefs.

(b) It could the case that in step (a) all but one set of preferences is eliminated. Then the agent is to adopt these as his overall preferences and form his strategy accordingly.

(c) If step (b) does not apply, the agent can treat the preferences that are not eliminated in step (a) as equally important. He may be able to avail himself to a method of aggregation or other selection procedure that determines the overall preferences he is justified to act upon. This method could be conceived of in analogy to making collective decisions where the preferences of several agents are inputs for the choice rule that determines collective preferences.

In Example 3-A, employing this procedure may result in (a) the agent deciding that his preferences at t_1 or at t_2 have no authority, in which case he should act upon the remaining one, or in (b) the agent deciding that none of has superior authority and finding a compromise between them.

Some comments are in order. First, there is no reason to think so far that if in problems with changing preferences there are chance moves, ex post justifiability would therewith be undermined. Second, there is no guarantee that the above procedure will deliver preferences that the agents always will have a strong desire to follow. But if it is accepted that the above procedure is part of rational deliberation, then the overall preferences, together with the other elements of the problem, will in principle provide all the reasons the agent has for deciding about an appropriate course of action. Third, once he sorted out his preferences according to the procedure, ex ante optimality conceived in terms of the overall preferences will derive criteria for ex post justifiability.

III

While in intertemporal decision problems with preference changes, it is usually assumed that preferences arising at future histories of the problem are determinate, clear foresight of one's future preferences might be exceptional. This does not always have to be the case. Consider then the next decision problem (Example 4). At the outset the agent still can choose between actions x and y, respectively. If he takes the latter, he will get 2 at t_3, no matter what. If he decides for x, he can make an other choice after t_2, between z and w. However, before that, at t_2, he will learn whether a possibility is realized, which is whether his preferences about the eventual outcome do or do not change. If θ holds, he will receive 4 if he does z, and 0 otherwise; and these payoffs also reflect his preferences at t_1. Now if the possibility is not realized, that is if ~θ holds, he attaches value 0 to doing z, and 1 to doing w. Suppose also that the agent is able to make an assessment (at t_1) of whether his current preferences will indeed change with probability p. Let me note that this example is an extension of a simpler one (Example 4*) in which the agent has no action to make after t_2, and where he assesses the eventual outcome, if θ, as having value 4, and if ~θ, as having value 1.

What should the agent do at t_1? Assume first that he will opt for z if the possibility realizes, and will choose w if it does not. Foreseeing this at t_1, and caring about his t_1 preferences, he will then choose x if p is larger than one-third. But why should he take his t_1 preferences so seriously? Indeed, we can recommend to him to employ the procedure outlined at the end of the previous section. Alternatively, if he regards his preferences that arise if ~θ as most authoritative, he should choose x only if p is smaller than one-third. But it is not entirely clear how that procedure could be applied to the analysis of the current kind of example. For instance, there may be no reason to regard preferences that are very unlikely to arise as even potentially authoritative (here p is either very large or very small). But one can think of a method that gives certain weights to potential preferences if they are to be aggregated. Let me propose to set this issue aside. Still, suppose that the agent chose x at t_1, upon havinf gone through the procedure, and that what he learns at t_2 is that the possibility did not realize. Should he regret his previous decision? Besides feeling obvious regret, should he think that his decision was not justified? We can state, for an answer, that in so far as he carried out the steps for determining which preferences are authoritative and which are not, and acts accordingly, he should be justified even if $\sim\theta$ is true.

Consider, however, the following case, to be referred to as the Problem Case. Suppose for this that the possibility that realizes at t_2 has to do with learning that the preferences the agent has at t_1 are really authoritative, that is that the agent learns something about himself, namely what his true preferences are. In the contrary case, when $\sim\theta$ holds, he learns that his preferences at t_1 are not really his. It follows that in this version of the problem he does not know at the outset what his true preferences are. So what should he do at t_1 in this case?

We are naturally interested in whether rational deliberation about what preferences he should act upon is possible here. Note that the decision problem of the Problem Case conforms to the model in section I. If the agent knows that he will learn his true preferences only at t_2, he cannot take his t_1 preferences as authoritative. But also, by the construction of this case, he cannot know which of the possible preferences he will have later are authoritative, so he cannot eliminate either of them. The aggregation of these potentially authoritative preferences is foreclosed, too. There is no advantage in thinking that the two preferences that depend on the possibilities θ and $\sim\theta$, respectively, taken together, are to trump his preferences at t_1. Now we might find reasonable to supplement the procedure for determining overall preferences with this further rule:

(d) If each preferences are eliminated in step (a), then the agent could just act on the preferences he has at the outset.

In the problem Case then a rational agent might do no more than to act upon his actual preferences, and choose x with an eye of continuing with z. We may consider the further proviso according to which x is to be chosen only if the probability of θ should be neither too large nor too small. But it is not clear, again, how probabilities can interact with other considerations in this case, neither how the range used in the proviso could be settled.

So suppose that the agent, having chosen x, indeed learns that his ex ante preferences were not authoritative, that is that $\sim\theta$ is the case. Can he think ex post that his decision was justified? His strategy is time-consistent if it calls for choosing z after θ and w after $\sim\theta$. If p is larger than one-third, or just large enough, his t_1 choice could be taken as ex ante optimal as well, in terms of his t_1 preferences. Further, he can assess retrospectively that he followed a deliberation procedure that was rational. Nevertheless, his decision for x cannot be taken as justified as it was not grounded on the authoritative preferences he knows he has. He has to acknowledge that given what he knows at t_2, his deliberation at t_1 did lack a sufficient basis. So we can say that the agent cannot satisfy himself with the thought that he decided well. We may even say that it is the outcome of the decision that will justify him. And the relevant outcome here, which is being able to carry out a plan generated by the preferences the agent has at the outset, depends on a contingency outside of his control.

One may then argue that the current case is not altogether different from standard cases of preference change. But in those cases, the agent could in principle know the status or the authoritativeness of his preferences at the various histories of the decision problem. Then his decision could indeed be justified. If he still does not know them, his decision will never be unjustified.

Is the Problem Case important? In many sorts of generic decision situations such a case will arguably not arise, or if it will, it will not give rise to much concern. One may be well aware that only going to a certain movie or to a certain shop will reveal whether one really cares about that film or that store, and still go; possibly being in fact even thankful if he eventually learns something new about his preferences once the program is over. In other cases, one just knows at the outset what one really wants and so the possibility of learning the contrary can just be discounted. We may say that it is in cases of decisions the outcomes of which come about only after a longer time period and when relatively much is at stake where belated learning could become a significant issue. Relevant situations include the choice of a hobby, or of a university to attend. Here it is plausible both that one does not know his true or authoritative preferences while making a decision and that the outcome is a matter of accumulated experiences. And what is at stake is whether one can identify with a course of action adopted at the phase of deliberation after learning much more what course really suits him.[2]

However, one could object that in these latter cases an agent may just not be aware that his preferences at the time of the deliberation are not genuine or authoritative, nor that he will at a definite point of time come to know whether this is the case. And these were crucial features of the Problem Case described above. I would like to remind, however, that the case is presented in confines of decision theory, where unforeseen contingencies have to be inscribed into the decision problem conforming to a canonical model. And the agent may well be aware that his ability to carry out a long-term project will depend on a possibility that has not unequivocally declared itself, on his capacity to identify with the project he prefers when he launches it. This awareness could be appropriately captured by the Problem Case.

[2] See Bernard Williams, "Moral Luck", in: Moral Luck (Cambridge University Press, 1981), pp. 31-36.

IV

When preferences change within a decision problem, the decision-maker is presented with reasons for acting in that problem that were not those that he had at the outset. But this is not the only way in which the stability reasons for acting in a problem could be upset, as acquiring or losing information in it may also give rise to new considerations for what to do.

Consider first the case of acquiring new information. Recall Example 2, where the agent's payoffs, if he chooses x, depend on the value of the variable θ. Suppose now that just as he is to actually make x, he learns that ~θ holds. Then he should avoid x, even if the probability of θ is larger than one-half, and thus the ex ante optimal strategy involved opting for x. Let this be Example 2*. This example does not fit into the canonical model of Section I, as the possibility of learning the value of θ ought to be foreseen before the decision problem unfolds. Now, on the other hand, the decision-maker could also lose information, as in the following Example 5. Here he can choose between x and y at t_1, and between z and w at t_2. If he chooses x and follows it up with z, he gets 4; If he chooses y and follows it up with w, he gets 2. Otherwise he gets 0. However, he will not remember at t_2 what he did at t_1, he cannot distinguish between the histories spanned by x and y, respectively. Nevertheless, that he will forget is foreseen in the ex ante state.

Let us ask the same question as in the context of changing preferences, is it possible in situations with forgetting that the ex ante optimal strategy, or plan, will not be justifiable ex post? The first observation to be made is that the agent may forget all sorts of things about the problem, while acting in it. Still, let us remain in the confines of the canonical model, and allow only those instances of imperfect recall where the information sets of the agent actually partition the histories when he is to make a move.[3] Next, it is to be acknowledged that for some of such problems with forgetting ex ante optimal plans may not be time-consistent. While this may be demonstrated for a variety of cases, I propose to focus on a well-known construction of Piccione and Rubinstein. In this example (Example 6), exhibiting absent-mindedness, the agent if he chooses y at t_1, will get 0. If he chooses x (at t_1), and then follows it up with x at t_2, he will get 4 at t_3; while if he takes y at t_2, will receive 1 at t_3. However, he does not know whether he

[3] The definition of imperfect recall problems is cumbersome; see Rubinstein, Modeling Bounded Rationality, pp. 67-70. However, the three examples below, Examples 5 to 7, are straightforwardly decision problems with forgetting.

is at history ∅, or at the history spanned by having opted for x at t_1. The ex ante optimal strategy is to do x. At the same time, when he is in his information set, he should think that the probability of being at t_1 is one-half. Then choosing y promises an expected payoff of 2, therefore he should do that.

As I argued in Sections II and III, time inconsistency does not settle whether the adoption of a strategy at the outset is justified ex post or not. The ex ante optimal strategy should reflect all the relevant reasons there are at the time when the agent still deliberates. Surely, the fact that information will be lost within the problem is one such reason. But I also claimed that the authoritativeness of preferences at the various histories ought also be taken into account, noting that raising this issue is legitimate in a theory of rational choice even if it cannot be handled within the canonical model. In problems with imperfect recall, an analogous difficulty arises with the interpretation of what a strategy is. While the standard definition of a strategy is that it renders to each information set of the decision-maker a certain action, or a probability distribution over the available actions; a strategy is simply a means of carrying out the course of actions that is regarded as the best in the ex ante state. In problems with preference changes it is an issue whether this strategy will actually control the behavior of the agent within the problem. But consider again Example 5. Here the ex ante optimal strategy calls for doing x both at t_1 and at t_2. At t_2, the agent will not know, by construction of the example and by the definition of information sets, whether he did x previously or not. Nevertheless he can remember his ex ante optimal strategy and just do x, and this will assure him payoffs at the level of 4. Does this mean that his information set at t_2 was not singleton after all? No, as he could just infer then from his strategy where he was at the problem and the definition of information sets does not restrict such inferences. Of course, it could be the case that an agent cannot recall his ex ante strategy. But this is an issue that can be represented only as a restriction on the set of available strategies, and there is no reason to impose this restriction always. Accordingly, I will assume that strategies can control the behavior of the decision-maker while in a decision problem, and part of this capacity is the capacity of their being remembered.

Let us turn now to the following example (Example 7). At t_0, chance decides whether the problem goes Left (L) or Right (R), with equal probability. Then at t_1, no matter what happened before, the agent can either choose x or y. If he chooses y, he gets 2. If he chooses x, he will have to make at t_2 a further choice between z and w. After R, the choice of z brings 4, whereas after L it brings 0. And after R, the choice of w brings 0, whereas after L, it brings 3. However, while the agent learns at t_1 the resolution of the chance move; he will forget this at t_2 (if he opted for x). Here the ex ante optimal strategy is to choose y after L, to choose x after R, and to choose z at t_2. But after L, the agent should wish to change this strategy; and opt for x followed by w. The issue is not simply that of the time inconsistency of the ex ante optimal strategy. Rather, whether it can be remembered at t_2 that the strategy was changed at t_1. Again, there is no principled reason to impose such a restriction; the fact that if the updating of strategy is recalled then the decision-maker could infer where he is at t_2 does not yield such a principled reason. I take it that the availability of such updating is a possibility that is analogous to the availability of judgments on the authoritativeness of history-indexed preferences in problems with preference changes: there is conceptual room for both of these within the framework of the canonical model.

The ex ante optimal strategies for the three imperfect recall problems presented in Examples 5 to 7 might call for updating, and can assure that the agent can reach the highest possible payoffs he could get in the associated perfect recall problems. In Example 5, he should start with x and continue with z. In Example 6, he should start with x, and after that update his strategy to prescribe y for the only information set he has to make a choice in. Finally, in Example 7, he should choose x after both L and R and z at t_2, and should update this strategy after L to prescribe w at t_2. Having carried out these strategies will then also be justified ex post.

I leave open the question here whether we can construct a decision problem with imperfect recall for which there is no strategy that will be regarded as justified ex post. This issue turns on the possibility of finding a reason for acting in a problem that is not available ex ante. But in the three problems with forgetting there are no such reasons.