

Strategic Rationality

Wolfgang Spohn
Fachgruppe Philosophie
Universität Konstanz
D-78457 Konstanz

Abstract

The paper argues that the standard decision theoretic account of strategies and their rationality or optimality is much too narrow, that strategies should rather condition future action to future decision situations, that practical deliberation must therefore essentially rely on a relation of superiority and inferiority between possible future decision situations, that all this allows to substantially broaden the theory of practical rationality, that a long list of points attended to in the literature can be subsumed under the broadened perspective (including a novel view on the iterated prisoner's dilemma and on iterated Newcomb's problem), and that the task to complete and systematize this list indeed forms a future research programme.

Content

1. The basic idea	1
2. The standard account of strategies	4
3. The general theory	10
4. Favorable and unfavorable changes of decision situations	18
5. Applying the model to iterated Prisoner's Dilemma and to iterated Newcomb's Problem	34
6. A list of further issues	43
Bibliography	52

1. The basic idea

Standardly, a strategy is conceived as an action plan which makes future action contingent on the given external situation. For instance, a complete chess strategy for White would have to specify a move of White for each possible position on the board or, if that should make a difference, for each possible previous course of the play. It seems obvious that this standard conception is much too narrow; a strategy should make future action contingent not on the future external situation, but on the

future decision situation of the agent – where a decision situation is something internal to the agent insofar as it consists of her beliefs and views, her aims and desires, or other determinants of rational action. The standard view is a special case of the broader one, since it includes the standard assumption that the agent gets informed about the external situation on which her strategy depends (otherwise it would be hard to understand how she could make her choice dependent on the external situation); thus one can as well view a standard strategy as making future action contingent on changed information states of the agent, thereby subsuming the standard notion under the broader view.

The minor point I would like to bring out in this paper is that the broader view is indeed essentially broader. There is, not surprisingly, a great variety of changes in decision situations. The doxastic side (concerning beliefs) may be affected in various ways: The agent may receive information, as just envisaged; experience may, however, change her doxastic state also in other ways. The agent may simply forget about things. There are various means, drugs for instance, for changing her outlook directly, i.e. without mediation of the senses. The agent may acquire new concepts and thereby restructure her decision situation. And so forth. The motivational side (concerning desires) may also be affected in various ways: Belief changes may induce changes in desires (insofar they are extrinsic ones about means and not intrinsic ones about ends). And intrinsic desires may and do change as well, for very diverse reasons: because of genetic disposition (we are simply born to have different desires at different ages), because of boredom (old desires get worn out), because of curiosity (new desires are discovered), because of education (wine tasting considerably forms the wine preferences), because of stupefaction (how much TV do you look?), again because of drugs (but there are other addictions as well), advertisement, brain-washing, and so forth.

Now, the minor point is not only that all these changes actually occur, but rather that we can, should, and actually often do take a strategic attitude towards them. We await new information – that was the standard case. We put on glasses in order to see sharper and thus to make more determinate observations. We write down things before we forget them. We usually use and do not reject occasions to receive education (which may, but need not consist in information). The weaned smoker abstains from the next cigarette, not because he could not like it, but because of fear to get addicted again. Ulysses ordered to tie him at the mast so that he could

not succumb to the chants of the Sirens. And so forth. I shall consider all these cases and many more in greater detail in section 4.

The major problem, however, is: What is the strategic rule governing such practices? What is the optimality criterion for strategies within this broader perspective? The problem is that the standard account cannot be carried over to the general case considered here. I want to propose a solution of this problem. Its basic idea is simple: The agent has to evaluate beforehand which of her possible future decision situation, or to put it more grandly, which of her possible future selves embody a superior, an inferior, or an incomparable point of view. For instance, getting informed or getting educated moves one to a superior position, whereas forgetting or getting brain-washed moves one to an inferior position; the position gained by a change of taste may or may not be incomparable; and so forth. The strategic rule then has to respond to the various kinds of positions; very roughly put, one better guards oneself against moves to inferior positions, does not care about incomparable ones, and seeks superior ones.

This sounds well enough, I guess. First, however, it looks a bit crude: Is the agent not required to have a finer evaluation of her possible future selves? I shall argue in section 6, (4), that this is not the case. Secondly, and more importantly, the basic idea so far sounds rather metaphorical. How could it be given a precise explication? I shall try to explain how this can be done. Clearly, this then allows for a general theory of strategic rationality.

Equally clearly, however, this general theory remains quite formal unless more is said about how possible decision situations are related as superior and inferior; it is this issue which gives substance to the theory. I shall consider quite a number of cases, in order to exemplify the general theory, in order to show its rich connections to the existing literature, and in order to give this issue at least a partial treatment. Indeed, the major point I would like to bring out in this paper is that this issue is a deep and consequential one well deserving more intense scrutiny.

So, the plan is this: In the next section I briefly review the standard account and state it in a simple way which is at the same time suited for generalization. Then, in section 3, the general theory of strategic rationality must be spelled out. Section 4 will substantiate the central notion of superiority by looking more closely to various applications. A further section is reserved for a specific and more tentative application to the iterated prisoner's dilemma. The last section 6 will briefly address a mixed bag of issues possibly deepening the understanding of the theory.

2. The standard account of strategies

We cannot proceed without referring to some standard formalization of decision situations. Here I shall take decision theory simply for granted, that is, I shall represent opinions, beliefs, or doxastic states in general as subjective probabilities and aims, desires, or motivational states in general as utilities. There are reasons for quarreling with these representations, but not here. Within these confines there are various variants of formalization. I comply here with the one given by Fishburn (1964) which still seems to me to be the most adequate and which runs as follows:

The first ingredient is a set V of variables; I call V the (*conceptual*) *frame* of the decision situation. V divides into a set A of *action* variables and a set X of other variables; let's call them *occurrence* variables (in default of an established name); thus $V = A \cup X$. What is a variable? Usually, it is a function from some state space left implicit into some set of values. It is simpler, however, to conceive of it just as a set of (at least two) values; the state space may then be taken as the Cartesian product $\prod V$ of all the variables in V , and the variables in the former sense then are the projection functions from that Cartesian product to its coordinates. The values of the action variables are to be chosen by the agent, simply by acting in this or that way, whereas the values of the occurrence variables realize in some way or other which may, but need not be influenced by the agent.¹

If V_1 is some variable in V , I shall use the corresponding lower case symbols v_1, v_1' , etc. to denote values (i.e. members) of that variable. And similarly, if W is some set of variables in V , I shall use the corresponding lower case symbols w, w' to denote the members of $\prod W$, i.e. the tuples consisting of values of the variables in W .

The precise nature of the values of the variables is unknown. Are they events or possible events? Are they states of affairs? Are they propositions suited as objects of attitudes like belief and desire? But what on earth are propositions? I am

¹ Here is the point where Fishburn (1964) deviates from Savage (1954). Savage distinguishes variables representing states of the world and variables representing consequences and then conceives of actions as deterministically relating states of the world and consequences. This narrow view of actions creates trouble. Fishburn avoids the trouble by not making this distinction and by assuming a primitive set of action variables. However, he generally speaks of consequence variables. This is misleading, because not every consequence variable in Fishburn's sense has to represent consequences of actions in the intuitive sense. My term "occurrence variable" is more neutral in this respect. For further comparison see Spohn (1977) and (1978, ch. 2).

convinced that decision theorists can learn a lot by looking at the obscure philosophical debates about these issues²; but again this is not our present concern.

I shall assume the set V of variables and each variable in V to be finite. The mathematical and conceptual difficulties emerging in the infinite case are intriguing, but I do not see that they are relevant to our present purpose.

The set V of variables has a temporal structure; the variables realize at different times. This structure is of utmost importance, also for this paper. Still, we need not burden our notation with it and may leave it largely implicit. For the present section it suffices to assume that the set A consists of the action variables A_0, \dots, A_m , to be realized at times $t_0 < \dots < t_m$, and that t_0 , the time of A_0 , is the time at or just before which the agent is situated; that is, the agent has presently to choose some action a_0 from A_0 , some time later some action a_1 from A_1 , and so on until she has chosen a whole course of action a from A . A_m forms the action horizon of her decision situation as formalized by the given frame V insofar as later actions are not considered in that frame.

A probability measure *for* a set W of variables is just a probability measure on the power set of $\prod W$. So, we next represent the doxastic state of the agent by a family $P = (P_a)_{a \in \prod A}$ of probability measures for V which is indexed by the courses of actions and for which $P_a(a) = 1$ for each $a \in \prod A$.³ The idea is that each P_a represents the agent's assessment of how likely the various events in the world are *given* that she takes the course of action a ; this also explains the condition $P_a(a) = 1$.⁴ I propose to add the further condition that any occurrence variable is probabilistically independent of all later action variables; i.e., if Y_k denotes the set of all occurrence variables before A_k , then $P_a(y_k) = P_{a'}(y_k)$ for all action courses a, a' which agree on A_0, \dots, A_{k-1} and differ at most in A_k, \dots, A_m . Since it seems correct to me to interpret subjective probabilistic independence in this case as subjective causal independence, I shall say that P , by satisfying this condition, is *agent-causally well-behaved*.

² Cf., e.g., Lewis (1979) or Spohn (1997) or (1998).

³ I shall always be so sloppy not to distinguish between, e.g., the course of action a and the set of the tuples v in $\prod V$ extending a ; the difference is purely formal and without intuitive counterpart. Alternatively, each P_a may be taken as a probability measure for X ; my variant will prove to be slightly more convenient.

⁴ The only, though crucial difference between Jeffrey (1965) and Fishburn (1964) is that Jeffrey additionally assumes the agent to have a probabilistic assessment of her own courses of action which is just what we need to add to the family P in order to get a probability measure for the whole of V . In Spohn (1977, section 5) I criticized this additional assumption, and Jeffrey's response (1977) left me unconvinced. I shall briefly return to this issue in remark (2) of section 6.

A utility function *for* a set W of variables is just a function from $\prod W$ into the set of real numbers. So, finally, the agent's motivational state is represented by a utility function U for V .⁵ U is not to represent any expected utilities within the given decision situation, whence I often refer to U as the agent's intrinsic utility function. However, it is important to recognize that intrinsicity as used here is not an absolute, but a relative notion. It is relative to the given frame V , insofar as the utilities specified by U are intrinsic only within V , but may well be expected utilities with respect to variables not explicitly represented in V .

Now a *naïve decision situation* δ at t_0 is just any quadruple $\langle A, X, P, U \rangle$ with A , X , P , and U as specified. I call such a decision situation naïve for the following reason: What is the rational or optimal thing to do in the decision situation $\delta = \langle A, X, P, U \rangle$? If I were naïve, I would think I had to choose a whole course of action in an optimal way, and I would apply the standard criterion of maximizing (conditional) expected utility. That is, I would define the *expected utility* of each course of action a as

$$EU(a) = \sum_{v \in \prod V} U(v) P_a(v)$$

and would opt for an a for which $EU(a)$ is maximal. If I were somehow forced to fix a whole course of action, this would indeed be the rational thing to do.⁶ Usually, however, I am not so forced, and then that procedure would be naïve because it neglects the fact that when the times t_1, \dots, t_m have come to choose from A_1, \dots, A_m I may be in different and hopefully better decision situations; it is only A_0 from which I have to choose immediately.

The essence of the strategic view of decision situations, which is the topic of this paper, precisely consists in taking this fact into account. It is clear that this requires a somewhat richer representation of decision situations, and the present aim is to find that enrichment which is most suited for our further discussion. The best thing to do for this purpose is to look at the standard account of the strategic view,

⁵ This is only more convenient, not more general than Fishburn's account. Fishburn assumes utility functions for X , but acknowledges that it may be required to have the basic utility function defined for A as well. In that case he simply, though inelegantly duplicates the set A within X ; cf. Fishburn, 1964, pp. 25f.

⁶ Of course, I am here neglecting the rich discussion about the justification of the Bayesian principle of maximizing expected utility which is, once more, not my present concern.

though this will be a formal exercise somewhat tediously extending over several paragraphs:

The standard account proceeds from the assumption that the agent may observe or experience certain variables, possibly at certain costs, and makes her decisions depend on the results of these observations. Thus, some of the action variables are *observation* variables; some of the occurrence variables are *observable*; with respect to an observation variable the agent has to decide whether to observe the relevant observable variable or not; and if she opts for observation, the observable variable will be *observed*. Moreover, all action variables are observed ones. They are observed, so to speak, by proprioception; however, their observation is automatic, involuntary and does not require a further *act* of observation; otherwise we would get into a vicious regress.⁷ We shall not need to make all this more explicit; let us just say, for any strategy s , for any time t_i , and for any course of events x from $\prod X$, that $o_{i,s,x}$ denotes the observed course of events before t_i as it is according to s and x . A *strategy* s then is a function which, for each i and x , assigns to $o_{i,s,x}$ a value from A_i .⁸ This is only a semi-formal way to say that a strategy makes each action dependent on what has been observed so far. In the end, each strategy s and each course of events x from $\prod X$ uniquely determines some course of action from $\prod A$ denoted by s_x , though possibly a different one for different courses of events from $\prod X$; this is precisely the point of strategies.

Given this description, the above optimality criterion for courses of actions is generalized to strategies in a straightforward way. The (conditional) *expected utility* of a strategy s is simply defined by

$$EU(s) = \sum_{x \in \prod X} U(s_x, x) P_{s_x}(x),$$

and the rational thing to do is to act according to a strategy with maximal expected utility.

However, it is quite obvious that this description of the strategic situation is not suited for further generalization to the cases envisaged above. The reason is that in these cases s_x and hence the given definition of the expected utility of strategies

⁷ This description of the active and the passive part of observation is very rough. It would be interesting to make it more adequate; but this need not concern us here.

⁸ This sounds circular, but, of course, it is just a brief simultaneously recursive definition. What I have omitted is the precise recursion clause defining $o_{i,s,x}$.

become meaningless; it is no longer true then that a strategy and a realization of the occurrence variables uniquely determine some course of action, since in these cases a strategy does not simply respond to observed external occurrences, but to possible decision situations coming about for various reasons which need not consist in observations of occurrence variables in the given frame. So, what to do?

There is a different, though equivalent description of the envisaged strategic situation. Let us focus on the backward induction step. That is, let us assume that we already have an adequate description of the possible strategic situations at time t_1 which says in particular which strategies s' there are at t_1 , which expected utilities $EU(s')$ they have, and thus which of them are optimal, and let us look at the step from t_0 to t_1 . We have to distinguish two cases:

The simple case is that the action variable A_0 is not an observation variable. Then the agent has simply to choose some option a_0 from A_0 and thereby gets into a decision situation at t_1 the only change of which, as compared with t_0 , is in the probabilities: the probabilities at t_1 are the ones from t_0 conditionalized by a_0 .⁹ Let us say in this case that the situation at t_1 *comes from* the one at t_0 *by action*. A strategy s at t_0 is then simply a combination of a choice from A_0 and a strategy s' at t_1 , and its expected utility $EU(s)$ is simply the expected utility $EU(s')$ of the strategy s' it implements.

The slightly more complicated case is that the action variable A_0 is an observation variable. The agent may decide not to observe the relevant observable variable, say X_0 . Then the case is as before. Or she may decide to observe X_0 . Then, again, the only change in the decision situation at t_1 she gets into lies in the probabilities: the probabilities at t_1 are the ones from t_0 conditionalized by x_0 , the actually observed value of X_0 , and by a_0 , the action to observe X_0 .¹⁰ Let us say then that the situation at t_1 *comes from* the one at t_0 *by (determinate) observation*. In this observational case, a strategy s at t_0 is a function which assigns to each observed value x_0 a strategy $s(x_0)$ for the resulting decision situation at t_1 , and its expected utility is defined as the expectation of the expected utility of the resulting strategies, i.e.

⁹ More precisely, the doxastic state at t_1 is then characterized by the family (P_a) for all action courses $a \in \Pi A$ starting with a_0 .

¹⁰ More precisely, the doxastic state at t_1 is then characterized by the family $(P_a(.|x_0))$ for all action courses $a \in \Pi A$ starting with a_0 .

$$EU(s) = \sum_{x_0 \in X_0} EU(s(x_0)) \cdot P_{a_0}(x_0).^{11}$$

This exhausts the description of the induction step which, however, should be put a bit more succinctly. To this end let us define that $\delta_0 = \langle A, X, \Delta, P, U \rangle$ is an *observational decision situation at t_0* if and only if $A = \{A_0, \dots, A_m\}$ and X are as before, Δ is a new occurrence variable, namely a set of (possible) observational decision situations at t_1 which come from δ_0 by action or by determinate observation, $V = A \cup X$ as before, $P = (P_a)_{a \in \Pi A}$ is, unlike before, a family of probability measures for $V \cup \{\Delta\}$, and U is a utility function for V and thus differs from P by not covering Δ .¹² This restriction should look intuitively plausible, and I defer a fuller defense to section 6, (5). A *strategy s for δ_0* then is a function defined on $\{A_0\} \cup \Delta$ which assigns a choice $s(A_0) \in A_0$ to A_0 and a strategy $s(\delta)$ for δ to each $\delta \in \Delta$. The *expected utility of s* , finally, is defined as

$$EU(s) = \sum_{\delta \in \Delta} EU_{\delta}(s(\delta)) \cdot P_{s(A_0)}(\delta)^{13},$$

where EU_{δ} is the expected utility function for the strategies for δ .

These are correct recursive definitions since backward induction must have a start at t_m or A_m , the action horizon of the decision situations according to the given frame; an observational decision situation at t_m reduces to a naïve one because Δ then degenerates to the empty set. It should also be clear that the two descriptions of the strategic standard situation are strictly equivalent. In a way, however, the second description just given is simpler than the first one insofar as it does not envisage the whole time span, but allows us to focus on one step in time by assuming that the rest of the time has already been treated adequately. Finally it should be apparent that the second description is indeed, unlike the first, suitable for generalization. So let us turn to the general theory the basic idea of which I have already indicated in section 1.

¹¹ Since I assume that P is agent-causally well-behaved, the probability of x_0 depends only on a_0 and not on later actions, and thus the last term becomes meaningful.

¹² This feature of U will be preserved throughout the definitions in the next section. It may seem unnecessarily restrictive. However, in my remark (7) of section 6 I shall defend this restriction.

¹³ Here the same remark as in the footnote 11 applies.

3. The general theory

Let us continue focusing on the backward induction step from t_0 to t_1 . Obviously, the only, though consequential point in which we have to generalize on observational decision situations consists in dropping the assumption that the possible decision situations at t_1 come from the earlier one at t_0 by action or by observation; in the general case now to be considered the change may result from any cause whatsoever: aging, forgetfulness, education, etc. This entails that we can continue using the conceptual apparatus developed so far. Let us therefore conceive a *general* decision situation δ_0 at t_0 as a quintuple $\langle A, X, \Delta, P, U \rangle$ as before, with the only difference that Δ is now a set of arbitrary *general* decision situations at t_1 which come from δ_0 by any way whatsoever so that nothing can be generally said about the relation of the probabilities and utilities at t_0 and at t_1 . So, the definition of *strategies for* δ_0 given above applies here as well. Again these are correct recursive definitions.

The crucial question, however, concerns the general optimality criterion for strategies. We can define the expected utility of the strategies for a general decision situation in the same way as above for an observational one and use the maximization of expected utility as an optimality criterion. In this case let us call the general decision situation a *positive* one, for reasons emerging later on. Thereby, however, the crucial question is only restated: Should general decision situations always be conceived as positive ones? Is maximization of expected utility in the sense defined generally reasonable? The answer is clearly no. Let me explain:

First, it should be clear that the agent, once she has moved to one of the decision situations at t_1 in Δ , will follow a strategy which is then rational or optimal; and as a strategically thinking person she has precisely this prediction about her future actions. Hence, whatever the appropriate criterion, it is constrained by the condition that it can recommend only those strategies at t_0 which result in strategies considered at t_1 to be optimal. Let us call this the *consistency constraint*¹⁴ which I

¹⁴ Introduced in the path-breaking paper by Strotz (1955/56) and scrutinized by Pollak (1968). Strotz discusses two strategies to face the threat of inconsistency, the strategy of precommitment and the strategy of consistent planning (which conforms to the consistency constraint). However, the first is only a special case of the second; it just consists in forcing oneself into a decision situation in which the actions which would be inconsistent in the situation arising without force are consistent.

take to be obligatory and beyond dispute.¹⁵

If we follow the optimality criterion for positive decision situations, we observe this constraint: we assess the possible strategies at t_1 according to their expected utility, and assuming that we shall maximize expected utility at t_1 , we maximize at t_0 the expectation of the possible maximal expected utilities at t_1 . At first sight this may generally seem a reasonable thing to do. It involves a hidden presupposition, however, namely that we now rely on our future assessment of the future situation; and this trust may very well be missing.

Let us look at a simple familiar example. Friends of mine living on the countryside have invited me to a dinner party, and I am thinking about how to go there. Public transport is not available; so I may either use my own car or take a cab. The evening will be entertaining and the wine delicious; so I foresee to get considerably, though still pleasantly drunk. If I came there with my own car, I shall have another decision problem at the end of the evening: How to get home? By driving by myself or by taking a cab? I know already what my assessment will be: taking a cab is expensive; the other day I annoyingly have to fetch my car somehow; the risk of getting controlled by the police is zero; and the risk of having an accident is as low as ever. So the expected utility of driving by my own will be much larger than that of taking a cab; driving by my own will then clearly be the rational thing to do.

If we were to stick to the optimality criterion for positive decision situations, the consideration would have to go on like this: If I now take a cab for getting to my friends, this costs 20 \$, and at the end of the evening I have no choice but to return by cab again. If, however, I take my own car, this costs almost nothing, and I shall have the pleasant choice of also returning with my own car. So, there is no question, to drive by my own to my friends is the rational thing to do.

Clearly, what I actually think is contrary to this criterion; it is this: I know that I shall be in a light mood after having drunk so much and that my assessment will be as described. But this is silly. Now, being sober, I judge that the chance of getting caught by the police is not zero, that the chance of having an accident is heavily increased for drunken drivers, and that the avoidance of these risks is well worth the prize of a cab back and forth. It is this assessment, and not the future silly one, on which I should base my present decision; so I better order the cab right now. I think the reasonableness of this consideration is immediately evident.

¹⁵ As does, for instance, Yaari (1977). The constraint has, however, been criticized by McClennen (1990) in a penetrating way. See my remark (3) in section 6 for further comment and comparison.

The deliberational structure in this case is opposite to the one in the observational case. In both cases I know what I shall do when I get into this or that decision situation; this is settled by the optimality criterion for strategies at t_1 which I have presupposed for the induction step and by the consistency constraint. The difference is this: In the case of observation I use my future evaluations in order to now assess how desirable it is to get into this or that decision situation; I base my present decision on my possible future evaluations. In the case of drunkenness, by contrast, I do not do so, finding my future evaluations inferior; rather I reassess the future situation also from my superior present point of view and base my decision on this assessment.¹⁶

If we grant all this, we face a problem. We have two special cases with convincing, though opposing optimality criteria, each of which bars the other one from serving as a general model. This entails that positive decision situations do not yet provide a general analysis. And it makes clear that we need a *principle of reevaluation* telling us in a general and precise way which evaluations we should employ in assessing or reassessing our future decisions.

I would like to propose the following solution to the problem which is, in a way, the core of the paper (though it is the shell around the core which makes it live). The idea is very simple and plausibly sounding: it is that *the reevaluation of a future situation δ has to proceed from the relevant superior point of view* which I shall call *the positive counterpart of δ* (and which may be identical with δ in the limiting case). This requires three explanations and refinements: What means "superior"? What means "relevant"? And is the "the" legitimate?

First, superiority is a primitive notion of my account. I just assume that each decision situation contains a relation \gg of superiority among the possible and counterfactual future decision situations presently considered which is a partial ordering, i.e. irreflexive and transitive; conversely, $\delta \ll \delta'$ says that δ is inferior to δ' . Superiority also induces a relation \approx of *evaluational equivalence* in the familiar way: $\delta \approx \delta'$ iff for all decision situations δ^* $\delta \ll \delta^*$ iff $\delta' \ll \delta^*$ and $\delta \gg \delta^*$ iff $\delta' \gg \delta^*$. I assume that we have a good intuitive grasp of the notion of superiority: if more is observed in δ than in δ' , then $\delta \gg \delta'$; if δ is forgetful as compared with δ' , then $\delta \ll \delta'$; and so on. Of course, there are many unclear cases, and in any case judgments of

¹⁶ Pollak (1968) criticizes Strotz (1955/56) for not having applied this second kind of consideration in a sufficiently consequent way and builds on it his own definition of the "sophisticated optimum path" (p.203).

superiority express subjective assessments and not objective facts. I shall consider many specific cases in section 4. Since we shall often consider changes of decision situations, let us moreover define a change from δ to δ' to be *favorable* iff $\delta' \gg \delta$ and *unfavorable* iff $\delta' \ll \delta$.¹⁷

Secondly, if I reckon with moving from the present decision situation δ_0 to the future situation δ , which is the situation from which to reevaluate what I would do in δ ? Not any situation superior to δ ; that would not make any sense, because there are too many superior situations and many extremely far-fetched ones, e.g. one in which I were to realize one of various unattainable moral ideals. Thus, we have to introduce a restriction to the "relevant" superior situation. The reasonable restriction seems to me to consider only that superior situation which is feasible in the sense that I would reach it if, in going from δ_0 to δ , I could cancel all the unfavorable changes and use instead the opportunities for favorable changes; the result of this counterfactual consideration is what I call the positive counterpart of δ . For instance, if the way from δ_0 to δ is favorable all the time so that $\delta \gg \delta_0$, then we have the limiting case in which δ itself (and not any situation superior to δ) is the positive counterpart from which to evaluate δ . If the change from δ_0 to δ consists in forgetting something or in getting drunk or is unfavorable in some other way, then the positive counterpart from which to evaluate δ is the unchanged δ_0 itself. If no change takes place from δ_0 to δ , but I could have moved to the superior δ' instead, then δ' is the right situation for reevaluation. Favorable and unfavorable changes may occur at once, as the dinner party of my friends might exemplify: drinking good wine may make me epistemically irresponsible, but it may also make me a more pleasant and less restrained chap. Then the relevant superior situation is the one which embodies only the favorable changes and neglects the unfavorable ones. And so on; but let us again defer details to section 4.

These preliminary examples make clear that the positive counterpart is often a counterfactual situation which the agent does not get into and does not think to get into. Still it is not a remote counterfactual possibility, but one within the reach of the agent; she would have reached it, if she would have been so lucky to undergo only favorable and no unfavorable changes.

¹⁷ Indeed, I wonder which of the two twin notions is intuitively the more basic one, the superiority/inferiority relation between decision situations or the characterization of situational changes as favorable or unfavorable. I have no clear answer and simply stick to my present practice.

However, to take up the third point, is the situation relevant for the reevaluation thereby uniquely characterized? This would be nice, but it is not generally the case. Hypothetically cancelling the unfavorable changes seems to have a unique result, but the superior situation to which I could move, though I do not, is not unique. For instance, I might make an observation, but I contemplate not to do so and thus to arrive at situation δ (which, hence, is almost the same as the one I am in). Making the observation would definitely be favorable. However, there are many possible results of the observation, and hence there are many possible superior situations $\delta_1, \dots, \delta_r$ I could get into and not a unique one. From which point of view should I then reevaluate δ ? Well, if there is no unique superior point of view for reevaluation, the only alternative I see is that my reevaluation is a mixture of the superior points of view $\delta_1, \dots, \delta_r$, where the weights with which they enter into the mixture are just the probabilities p_1, \dots, p_r with which I think I could arrive at $\delta_1, \dots, \delta_r$. If we call this mixture $\mu = \{\langle \delta_1, p_1 \rangle, \dots, \langle \delta_r, p_r \rangle\}$ a *mixed (decision) situation*, the lesson is that the positive counterpart from which to reevaluate δ is generally to be conceived as such a mixed situation. It needs to be emphasized that every bit of such a mixed positive counterpart may be purely hypotheticalal.

Formally, however, the last consideration turns out now to be rather a detour: The pure situations $\delta_1, \dots, \delta_r$ mixed in μ do not enter into the above consideration in an individually relevant way (as they would, for instance, if one had to decide what to do if one were in one of them); they are needed only as parts of a mixed reevaluation. For this purpose, however, the mixed point of view μ can be integrated into a single situation δ' in the following way: If $\mu = \{\langle \delta_1, p_1 \rangle, \dots, \langle \delta_r, p_r \rangle\}$ is the positive counterpart of δ , and if $\delta_i = \langle A, X, \Delta, P_i, U_i \rangle$ ($i = 1, \dots, r$), we may define $P' = \sum_i p_i P_i$, $U' = \sum_i p_i U_i$ (for $i = 1, \dots, r$) and $\delta' = \langle A, X, \Delta, P', U' \rangle$. For evaluative purposes δ' is clearly equivalent to μ , since it does not matter at which stage of calculating expected utilities we mix with the p_i . The upshot is that we may identify the positive counterpart of δ with the pure situation δ' instead of the mixed situation μ . I shall do so in future. However, one must always be clear about how δ' is to be construed; it need not be a superior situation at which the agent might have arrived under the specified counterfactual condition, but may as well be constructed from a mixture of such situations.

Let me schematically summarize these somewhat abstract considerations. I am now, at t_0 , in the situation δ_0 , pondering which action a_0 to take from A_0 thereby arriving at $\delta_1, \dots, \delta_n$ with probabilities p_1, \dots, p_n . It is presupposed in the induction

step that my decision rule already tells me that s_1, \dots, s_n are evaluated, respectively, in $\delta_1, \dots, \delta_n$ as optimal strategies to proceed. The consistency constraint thus allows me to assume that I shall proceed according to s_i once I am in δ_i . However, I need not rely on these evaluations, and so I reevaluate s_1, \dots, s_n . I do this by associating with each δ_i the mixed situation $\mu_i = \{\langle \delta_{i1}, p_{i1} \rangle, \dots, \langle \delta_{ir_i}, p_{ir_i} \rangle\}$, or rather the corresponding pure situation δ_i' , as the positive counterpart of δ_i . As I have explained, this is to be understood in the following way: If instead of moving to δ_i I had avoided all unfavorable changes and taken all feasible favorable changes, I would have arrived at $\delta_{ij} \gg \delta_i$ with probability p_{ij} .

Then I can reevaluate s_i as follows: In the superior δ_{ij} s_i would have had the expected utility $EU_{ij}(s_i)$, and my reevaluation $EU_i'(s_i)$ of s_i from the point of view of the mixed μ_i or, equivalently, the pure δ_i' is the expectation of these values, i.e. $EU_i'(s_i) = \sum_j p_{ij} EU_{ij}(s_i)$ (for $j = 1, \dots, r_i$). This is my *principle of reevaluation*. So, if s_0 is the strategy of first doing a_0 and then s_i after getting into δ_i , I finally arrive at my evaluation of a_0 : it is $EU_0(a_0) = \sum_i p_i EU_i'(s_i)$ (for $i = 1, \dots, n$). And my *decision rule* or *optimality criterion* is to choose that consistent strategy which maximizes this evaluation among all consistent strategies.

This may seem complicated, but in all the cases we shall consider it boils down to something quite perspicuous. Let us look, as a first illustration, at the two examples we have considered so far:

The first example consisted in the observational decision situations. Their subsumption is straightforward: If the agent is in δ_0 and decides to observe, she arrives at one of several superior situations which is its own positive counterpart. And if she decides not to observe and thus moves to δ (which is just as δ_0 with the exception that the decision not to observe is already past), the positive counterpart of δ is just the mixture of the situations informed by the observation which is in fact identical to δ .¹⁸ Thus, the above decision rule reduces to the rule for positive decision situations, i.e. to maximizing the expected utility of the strategies.

Our second example was the boozy dinner party in which only unfavorable changes were envisaged. Our optimality criterion takes account of it in the following way: Let t_0 be the time to get to the dinner party, t_1 the time late in the night to get home, c_1 the action to take a cab home, d_1 the action to drive home by myself, and c_0 and d_0 the corresponding ways to get to the dinner party at t_0 . At t_0 I am in δ_0

¹⁸ This is characteristic of observational changes. See section 4, (2), for more on this.

having two options, c_0 and d_0 . If I take c_0 I am sure to get into a degenerate decision situation δ_c at t_1 leaving no choice but c_1 ; the course $\{c_0, c_1\}$ has expected utility -40, we may assume, viewed from δ_c as well as viewed from δ_c 's positive counterpart which embodies my judgment at t_0 . If I take d_0 I am sure to get into a decision situation δ_d at t_1 in which I can choose between c_1 and d_1 . Within δ_d the expected utility of c_1 is -30, say, and that of d_1 -5. However, within δ_d 's positive counterpart which embodies my judgment at t_0 the expected utility of c_1 is -30, again, but that of d_1 is -50.

How then does the above decision rule apply to this situation? Trivially, c_1 is optimal in δ_c , and d_1 is optimal in δ_d . So, there are only two strategies, s_c and s_d , continuing optimally: s_c says first to do c_0 and then to do c_1 in δ_c (and to do d_1 in δ_d), and s_d says first to do d_0 and then to do d_1 in δ_d (and to do c_1 in δ_c). But now we have to evaluate s_c and s_d in the positive counterpart of δ_c and δ_d . And there the expected utility of s_c is -40 and that of s_d -50. So, s_c turns out to be the only optimal strategy for δ_0 according to our definition.

In principle, the schematic description given above is good enough for dealing with a rich variety of changes of decision situations. So far, however, we have only considered the recursive step from t_0 to t_1 and not yet taken account of the fact that this kind of deliberation may repeat on every stage of the recursion. Hence the task to turn all this into a correct recursive definition is still before us, and I should show how this task may be achieved (however, you may as well skip this part):

First, we must be able to say how the correspondence between decision situations and their positive counterparts extends through the whole recursion. This purpose is served by the following auxiliary definition: Let $\delta_0 = \langle A, X, \Delta, P, U \rangle$ and $\delta_0' = \langle A, X, \Delta', P, U \rangle$ be two general decision situations (cf. the first paragraph of this section). Then $\langle f, g \rangle$ is an *isomorphism from δ_0 to δ_0'* if and only if f is a bijection from Δ onto Δ' and g a function defined on Δ such that for each $\delta \in \Delta$ $g(\delta)$ is an isomorphism from δ to $f(\delta)$. In this way the isomorphism recursively extends through all times up to the action horizon.

Such a correspondence also enables us to identify corresponding strategies: Let $\langle f, g \rangle$ be an isomorphism from δ_0 to δ_0' , s a strategy for δ_0 , and s' a strategy for δ_0' . Then s' *corresponds to s via $\langle f, g \rangle$* if and only if $s'(A_0) = s(A_0)$ and if for each $\delta \in \Delta$ $s'(f(\delta))$ corresponds to $s(\delta)$ via $g(\delta)$. Obviously, this correspondence is unique.

Thus equipped we can proceed to our central definition: $\delta_0 = \langle A, X, \Delta, P, U, f, g \rangle$ is a *fully strategic decision situation at t_0* if and only if A , X , P , and U are as in a general decision situation, if Δ is a set of fully strategic decision situations at t_1 , if f is a function defined on Δ , and if $\langle f, g \rangle$ an isomorphism from δ_0 to δ_0' , where $\delta_0' = \langle A, X, \Delta', P', U \rangle$ is a positive decision situation at t_0 with $\Delta' = \{f(\delta) \mid \delta \in \Delta\}$ and P' being the probability measure for $A \cup X \cup \{\Delta'\}$ corresponding to P .¹⁹ Here, for each $\delta \in \Delta$ $f(\delta)$ is just what we have called the positive counterpart of δ . The relation of superiority characterizing an agent in δ_0 is thus implicitly contained in the isomorphism $\langle f, g \rangle$.

Finally, the *decision rule* or *optimality criterion* associated with a fully strategic decision situation $\delta_0 = \langle A, X, \Delta, P, U, f, g \rangle$ at t_0 is this: A strategy s_0 for δ_0 is *optimal* if and only if for each $\delta \in \Delta$ $s_0(\delta)$ is an optimal strategy for δ , and if for all strategies s for δ_0 the following holds: if for each $\delta \in \Delta$ $s(\delta)$ is an optimal strategy for δ , if s_0' and s' correspond, respectively, to s_0 and s via $\langle f, g \rangle$, if δ_0' is the positive decision situation defined from δ_0 as above, and if EU its expected utility function, then $EU(s_0') \geq EU(s')$.

This definition expresses and summarizes the considerations above. On the one hand it satisfies the consistency constraint since only a strategy can be chosen which implements optimal substrategies in the decision situations actually arrived at. On the other hand it embodies my principle of reevaluation since all consistent strategies have to be judged not according to how their value is actually seen later on, but according to their value in their (usually hypothetical) positive counterpart.

The final statement of my general theory²⁰ has certainly become unspicuous. Indeed, so many recursions are nested in a fully strategic decision situation that I dare say that no person has ever conceived of her situation in a fully strategic way.²¹ We should not be bewitched, however, by this complexity. It finds its ju-

¹⁹ That is: if $\Delta = \emptyset$, then $P' = P$, and if $\Delta \neq \emptyset$, then $P'_a(v) = P_a(v)$ and $P'_a(f(\delta)) = P_a(\delta)$ for all $a \in \Pi A$, $v \in \Pi V$, and $\delta \in \Delta$. It should be clear, moreover, how to extend the definition of an isomorphism from general to fully strategic decision situations.

²⁰ The above definitions may be already found in Spohn (1978), pp.171f., in full technical detail. The are only two differences. For one, I did not yet see the formally inconsequential necessity to allow positive counterparts to be mixed decision situations. The more severe restriction was that I took (determinate or indeterminate - see section 4, (3a)) observation to be the only way to move to a superior situation. Now, however, I think that there are more ways and hence that the notion of superiority is much wider and more interesting.

²¹ If you like combinatorics: if a fully strategic decision situation considers m successive times of action (not $m+1$ as assumed in the definitions), if all positive counterparts involved are pure, i.e. unmixed situations, and if all the non-empty Δ 's of all the general decision situations of the form $\langle A, X, \Delta, P, U \rangle$ involved in the fully strategic situation have the same number of elements, namely

stification in the fact that many different cases can be, it seems, adequately subsumed under these general definitions. We should indeed look at all these cases; they are much simpler, and only they give substance to the so far very abstract theory.

4. Favorable and unfavorable changes of decision situations

Let us thus consider, in a more systematic way, specific kinds of changes of decision situations, whether they are favorable or unfavorable, and whether the strategic recommendations entailed by the general theory seem plausible. The changes I shall discuss fall into three groups, doxastic changes, conceptual changes, and motivational changes, each with various subgroups. Concerning the first group we have already discussed:

(1) *Changes by action*, in which the agent chooses some act a_0 from the non-observational action variable A_0 , is proprioceptically aware of her choice, and finds herself facing a choice from A_1 , knowing she has done a_0 . This is hardly any change at all leading to an evaluationally equivalent point of view which is identical to its positive counterpart, since no superior point of view is under consideration in this change. Hence, the new optimality criterion reduces to the standard theory.²²

(2) *Changes by determinate observation*, in which the agent chooses "observe X_0 " or "do not observe X_0 " from the observational variable A_0 , gets sure about the actual value of X_0 in the case of observation, and conditionalizes with respect to this value of X_0 . As already mentioned, this is the case standardly considered in the decision theoretic literature. How does our theory treat it? In general, the observation is favorable, since it moves the agent to a superior, because better informed situation. Hence, the situations resulting from observing are their own positive counterparts, and their mixture (weighted by the probabilities of the

$n \geq 2$, then $q = (m \cdot n^m - 1)/(n - 1) - (n^{m-1} - 1)/(n - 1)^2$ decision situations are nested in the fully strategic situation. Thus, if $m = 1$, then $q = 1$, as expected; but if, e.g., $m = 4$ and n is minimal, i.e. = 2, then we already have $q = 56$.

²² One may envisage the possibility that one does something one did not want to do and thus arrives at a situation which is inferior in some sense, but not in the sense that it would be in need of reevaluation from the present point of view.

possible results of the observation) is the positive counterpart of the situation arrived at by not observing. However, the mixture is identical with that situation, and therefore our theory reduces to the standard theory in this case, too.

I said that observation is generally a favorable change. This indicates that there may be exceptions, as the progression from subcase (2a) to (2c) is intended to show:

(2a) The simplest case is the one in which the observation is *cost-free*. In this case a nice theorem holds saying that making a cost-free observation is always at least as good as making no observation. Skyrms traces this theorem back to Frank Ramsey in the 1920s and generalizes it to various infinite cases and to the kinds of doxastic changes considered below in (3) and (5) (cf. Skyrms 1990, pp. 94ff. and ch. 4).

(2b) The case of *costly* observation is more intricate because the comparison of the observation costs and the utility of the information received may be quite difficult. Usually, it is just the making of the observation which causes costs. Within our framework this may be represented in two different, though equivalent ways. Either, the utility function V of the initial decision situation δ_0 gives a lower value to all courses of events in which the observation is made than to the corresponding courses of events in which the observation is not made. Or the utility functions of the possible decision situations arrived at through the observation are decreased as compared with the utility function of the decision situation resulting from making no observation (where the former situations are still superior to the latter despite this decrease). Obviously, the second treatment is the more flexible one.

(2c) So far, we did not encounter the possibility of changes by observation to be unfavorable or at least not favorable. However, there are various examples suggesting this possibility; I subsume them under the heading "*unwanted information*". Many parents like to discover the sex of their child only at the delivery and not earlier. There is a level of human interaction (which may be very close in some respects) at which one often does not want to have more intimate knowledge about the other ones (whether one is disliked by them, what their sexual preferences are, etc.). The diary of my good friend, though it may be utterly revealing, is taboo for me. An employer may not raise any information whatsoever about his employees. And one may reasonably argue about moral limits of scientific knowledge; perhaps we better had not known about nuclear fission and fusion and perhaps we better should not know how to manipulate the human genom.

May these kinds of examples be rationalized? Yes, in various ways. The break of the pleasant suspense caused by premature knowledge of the sex of the child or the remorse caused by my breaking my friend's privacy may be seen as general observation costs as envisaged in (2b). However, the observation costs may also depend on the piece of information one receives. If the doctor's expensive diagnosis concludes in my having a curable disease, that's tolerable; but if it concludes that I am doomed this causes extra costs because I would spend the rest of my life in a lighter mood if I would not know about my disease. If I learn that my collaborator's attitude towards me is not adverse, collaboration can go on as before; but if I learn that he hates me, though he has never shown it, I cannot continue with the previous ease, and hence I would incur extra costs with this piece of information. Such cases are different from (2b). Indeed, the first way of accounting for observation costs which I mentioned does not work for such cases, since what is relevant for them is not the negative utility of the fact one gets informed about, but the negative utility of one's being informed about that fact. However, the second way considered in (2b) is clearly generalizable to the cases at hand.

Is this rationalization enough? I am unsure. As long as the doxastic states arrived at in these cases count as superior, the standard decision rule makes it a matter of weighing whether it is rational to acquire a piece of information; this depends on whether or not the benefit of information outweighs its costs. Since this may not always seem appropriate, it is worth mentioning that our general theory admits of a more radical treatment of these cases. If one considers decision situations brought about by unwanted information as inferior, then the positive counterparts of these situations contain the prior uninformed probabilities and the prior utilities reduced by the observation costs; if these are not negligible, our general optimality criterion tells that it is never rational to move into such an inferior decision situation, i.e. to gather unwanted information thus conceived. This may be what one wants in such cases as my reading my friend's diary or the scientists' uncovering secrets which better remain secret; but I am not sure. However, the fact that our theory opens this kind of consideration is certainly welcome. If so, this is only a nice side result; we saw and will see that the examples forcing the consideration of superior and inferior situations are different ones.

(3) *Further informational changes*: The philosophical literature discusses various kinds of changes of subjective probabilities besides the simple conditionalization considered so far. I would like to briefly mention the most important ones.

(3a) The conditions for observing the observable X_0 may be poor in various ways and thus not allow to determine the actual value of X_0 with certainty. It may, for instance, be too dark for clearly seeing the colour or the shape of an object, the connection may be too noisy for clearly understanding my interlocutor, or the variable actually observed may be one outside the given frame and the determinate observation of that variable may permit only uncertain conclusions about X_0 . In all such cases of *indeterminate observation* the belief change is appropriately described by Jeffrey's generalized conditionalization.²³ To be specific: If P is the agent's probability measure for X at t_0 (let us forget for the moment about the indication by courses of action), if the indeterminate observation of X_0 induces her to have the probability measure Q for X_0 , and if P' is her probability measure for X at t_1 after the observation, then P' is given by

$$P'(h) = \sum_{x_0 \in X_0} P(h | x_0) Q(x_0) \text{ for any event } h \subseteq \Pi X.$$

(3b) In (2) and (3a) we have made a hidden assumption which it is time to make explicit. In (2) we have assumed that the actual value, whatever it is, of the observable X_0 will be observed, i.e., that the possible pieces of evidence are the possible values of X_0 and hence form a partition. The same was assumed in (3), even though no piece of evidence was there taken to be received with certainty. This assumption is plausibly satisfied in genuine observation. However, when information comes in a more indirect way such as in the third example of (3a), the case is less clear. Indeed, Wagner (1992) gives a number of examples of *non-exclusive information* in which the set E of possible pieces of (uncertain) evidence does not need to be a partition of the possibility space, and he convincingly argues that conditionalization should then take a slightly more general form, namely that there should be a non-negative function m on E with $\sum_{e \in E} m(e) = 1$ (which can hence not be understood as a probability function, if E is not a partition) such that

²³ Cf. Jeffrey (1965), ch.11. Indeed, it was this chapter which initiated the current discussion on the dynamics of subjective probabilities.

$$P'(h) = \sum_{e \in E} P(h | e) m(e) \text{ for any event } h \subseteq \prod X.^{24}$$

²⁴ For details see Wagner (1992); he particularly emphasises the fact that m is like a basic probability assignment underlying a Dempster-Shafer belief function.

(3c) Evidence may come in still other forms. There is a principle which is able to cope with any evidential constraint whatsoever, *the principle of minimizing relative entropy*, or MAXENT for short, which runs as follows: If P and P^* are two probability measures for X , then the entropy $H(P, P^*)$ of P^* relative to P is defined as

$$H(P, P^*) = \sum_{x \in \Pi X} P^*(x) \log(P^*(x)/P(x)).$$

Now evidence may be understood as imposing some constraint on the new probability measure P' to the effect that P' should belong to some set C of probability measures for X – for instance to the set of measures P^* such that $P^*(x_0) = 1$ for some $x_0 \in X_0$ (cf. (2)) or such that $P^*(x_0) = Q(x_0)$ for all $x_0 \in X_0$ and some given Q (cf. (3a)). Then MAXENT says that P' is to be that measure P^* for which $H(P, P^*)$ is minimal among all P^* in C . There are mild conditions, for instance that C be convex and closed, which guarantee the existence and uniqueness of P' thus determined (cf. Csiszar 1975).

There is an argument whether or not MAXENT may be understood as an update rule (see, e.g., Skyrms 1987 and Hunter 1991), which I shall not attempt to dissolve.²⁵ (Cf., however, my remarks in (5)). So I mention MAXENT here only as a further very general possibility of accounting for probabilistic belief change.

The important question for us is, of course, whether the changes thus described should be taken as favorable. The authors proposing and discussing the mentioned rules clearly did take them so, as is also shown by the examples they give. However, it is equally clear that it is not the mathematical form of the changes which makes them so; any wild and unfavorable belief change may accidentally (or even not so accidentally – see (4)) take on one of these forms. Rather such changes are favorable when and because we think that we thereby get informed about the state of the world, that these changes are not endogeneous changes somehow produced by the doxastic system itself, but exogeneous changes somehow induced by the external world; and the rules mentioned try to capture at least necessary conditions of a formal kind for such changes.

Still, together with the decision rule for positive decision situations, which is the only one we need for these cases, these rules of belief change have a consider-

²⁵ I should also mention that Wagner (1992), p. 255, shows his generalization (3b) over (3a) to be incompatible with MAXENT.

able explanatory, i.e. rationalizing power. I already mentioned the fact, that curiosity is rational, i.e. that cost-free observation, be it determinate as in (2a) or indeterminate as in (3a) is never bad; I don't know whether this result extends to the case where the changes I expect to undergo are favorable, cost-free, and of the form (3b) or (3c). But explanation extends further; clearly, more determinate observation is usually more and never less valuable than less determinate observation. Hence, unless the costs are excessive, we are well advised to make observation more determinate. Indeed we take great pains to achieve this. We put on glasses, we construct extremely accurate measurement devices, we teach our children to have a clear handwriting, we make traffic signs salient, etc. Again, it would be interesting to know whether such results can still be generalized, whether, for instance, exclusive information as in (3a) is better than non-exclusive information as in (3b) or how valuable the various kinds of constraints on the posterior probability measure are.

(4) There is a great variety of *unfavorable doxastic changes*. Certainly the most common form is forgetting. Sometimes dreams and hallucinations are mistaken for reality. Drugs like alcohol (see my example above), cocaine etc. may have direct or indirect effects on our subjective probabilities. There is wishful thinking, and there is the whole range from mild persuasions over massive make-beliefs up to brain-washing. And so forth.

There are also a number of ambiguous cases. There are various methods, for instance psychoanalysis, for raising recollections; these may well be considered as favorable. Hypnosis is a more dubious method which may be used for bad doxastic influences as well. There are experiences of illumination (if these can count as doxastic changes) which I personally would utterly mistrust, but others may hope for. And so forth, again.

Do any laws of rationality hold for such changes? This appears to be a strange question. However, it does not seem implausible to me that Jeffrey's generalized conditionalization (3a) applies to such cases as well. That is, if the change directly affects only a certain kind of propositions forming a partition – for instance, I have only forgotten and thus am uncertain at which specific times the bus leaves, but have kept, of course, all my general knowledge about public transport – Jeffrey's formula is arguably a rational guide for the other posterior probabilities.

Which recommendations does our general criterion of optimality give in the case one expects unfavorable doxastic changes to occur? In the most elementary

case the agent has a choice between an unfavorable change and no change at all. Hence, the unchanged situation is the positive counterpart of the unfavorably changed one. The strategy which is optimal in the latter gets therefore reevaluated from the point of view of the former and then turns out to be at most as good as, and usually worse than, the strategy which is optimal in the unchanged situation. So, finally, the agent is well advised to lower the probability of an unfavorable change as much as she can – provided this does not cause any costs. If there are costs, then it is a matter of calculation of whether or not they are worth the avoidance of the unfavorable change.

This simple case is symptomatic of more complex situations, and it is as it should be, I think. For instance, it explains our continuous fight against forgetting. There surely are many good reasons against forgetting, but the most straightforward is the practical one which is well captured by our criterion of optimality. Also, our mistrust against all influences of the above-mentioned unfavorable kinds is thereby rationalized. These are explanatory achievements of which I am unaware in the literature.

(5) There is a general principle guiding only acceptable changes, the *reflection principle* proposed by van Fraassen (1984). It says, slightly adapted to our framework, that, if P is the agent's subjective probability at t_0 , Δ the set of decision situations she envisages to face, and P_δ her subjective probability in $\delta \in \Delta$, then

$$P(h \mid P_\delta(h) = p) = p \text{ for all } h \subseteq \prod X,$$

i.e. the agent's present judgment given her future judgment conforms to that future judgment.²⁶ Hild (1998a) aptly speaks here of an auto-epistemic principle, since it requires the agent to reflect not only on the subject matter at hand, but also on her possible opinions about it. This reflexive attitude is indeed characteristic of the general approach pursued here.

The reflection principle is easily seen to be equivalent to what is called the *iteration principle*, namely that the agent's present opinion is a weighted mixture of her possible future opinions, the weights' being her present probabilities for arriving at the future opinions, i.e. formally, that

²⁶ For a more careful statement of the principle see Hild (1998a).

$$P'(h) = \sum_{\delta \in \Delta} P_{\delta}(h) P(\delta) \text{ for all } h \subseteq \prod X.^{27}$$

Indeed, it is this principle which is at the base of my claim that in the observational case the situation resulting from non-observation is identical with its positive counterpart which is a weighted mixture of the possible situations achieved through observation.

There has been an argument about the validity of these principles. Cases were presented which I take to be clear counter-examples.²⁸ It is equally clear, however, that all these counter-examples involve unfavorable doxastic changes. So Jeffrey (1988), p. 233, proposed²⁹ to restrict the principles to those cases in which the agent assumes that all the doxastic changes she envisages are rational or, in my terminology, favorable ones.³⁰ Given this restriction, Hild (1998b) gives a Dutch book argument in favor of these auto-epistemic principles.

This raises the question how these principles relate to the rules of doxastic change mentioned in (2) and (3). Hild (1998a,b) argues, convincingly in my view, that simple and generalized conditionalization (2, 3a), and hence also the Dutch book arguments given for them, rely on hidden assumptions, the most important being the assumption of evidential independence which says that given the observed fact the events $h \subseteq \prod X$ are probabilistically independent from believing in (or having a certain probability for) the observed fact. Given these assumptions the conditionalization rules are entailed by the reflexion principle. If these assumptions are violated, however, these rules fatally contradict the reflexion principle; Hild (1998a) gives various examples for such violations, among them such problems as Freund's puzzle or the puzzle of the three prisoners.³¹

²⁷ This condition is explicitly stated in Spohn (1978), pp. 161f. (Def. 29). See also Goldstein (1983). For variations, more careful statements, and the equivalence proof see again Hild (1998a).

²⁸ Cf., e.g., Maher (1992) and (1993), p.107ff.

²⁹ In fact, however, this restriction was recognized already in Spohn (1978), pp. 162 and 166.

³⁰ Van Fraassen (1995) tries to defend instead of revise the reflection principle. I remain unconvinced, however. In what he calls the death and disability defense (pp. 22ff.) van Fraassen suggests that it may not be *my* future opinion to which unfavorable changes lead; rather, my mental life has then ceased or paused, in a sense. However, I don't think that personal identity should be conceived so narrowly. The integrity defence (pp. 24ff.) appeals to the preservation of epistemic integrity as a presupposition of the reflection principle. This seems very close to me to Jeffrey's restriction. And I simply can't see that van Fraassen (on pp. 21f.) has avoided the contradiction between the reflection principle and foreseen memory loss.

³¹ Indeed, Hild's observations account for a worry concerning unrestricted application of conditionalization which was put forward by Weirich (1983).

The case (3b) of non-exclusive information may also be embedded in this auto-epistemic setting. In this case, the measure $m(e)$ cannot be understood as the posterior subjective probability of $e \in E$. However, $m(e)$ may be auto-epistemically interpreted as the probability of having received the evidence e , and then Wagner's generalized conditionalization (3b) again follows from the reflexion principle. Finally, I do not know how the MAXENT rule mentioned in (3c) precisely relates to the auto-epistemic principles. It should be mentioned, however, that MAXENT, after being attacked by Skyrms (1987) as an update rule, has been defended by Hunter (1991) just within the auto-epistemic setting. Hence, it seems, despite the remaining unclarity, that the reflexion principle is the most general principle applying to acceptable doxastic changes.

However, it should be clear again that it is only a necessary and not a sufficient condition for favorable changes. The reflexion principle may be accidentally satisfied also in other cases. Hence, it is not this satisfaction which *makes* the changes favorable.

(6) Let us turn now to *conceptual changes* about which I have to say the least though they may well not be the least important. How concepts and meanings – which should presumably be carefully distinguished – are best to be construed, is still not clear among philosophers of language and mind. Even more we may speculate about conceptual and meaning change. Interesting as this may be, we better look only very superficially at the most primitive case.

This is the case of conceptual refinement which may be modelled within our framework simply by letting a decision situation contain new variables or more fine-grained old variables; both additions represent new distinctions. If the change from a prior to a posterior decision situation consists only in a cost-free conceptual refinement, and if the new distinctions have only doxastic import, we should expect two things: first, the marginalization of the posterior probability measure to the old and less fine-grained variables should be just the prior probability measure; otherwise, the change must include also non-conceptual, empirical information about the old variables.³² Second, the posterior utility of a more detailed course of events should just be the prior utility of the less detailed course of events of which the former is an extension (to the new and more fine-grained variables of the posterior

³² I suppress here well-known philosophical worries about the separability of empirical and purely conceptual information.

situation); otherwise, we must assume that the new distinctions have cause costs or have also evaluative import.

How may such conceptual change strategically to be taken into account? Here I have to admit that my formal framework does not really allow to do this, simply because I have assumed that all decision situations nested in a fully strategic decision situation are based on the same conceptual frame V . It would obviously be not very difficult to raise to the required formal generality. I did not do this because my general theory is already unperspicuous enough. Moreover, there is something odd in strategically considering conceptual change. For instance, one cannot foresee a conceptual change, since one could do so only by already having the new concept. Similarly, it is difficult to envisage possible conceptual changes, since for doing so specifically one should already know the possible new concepts. Hence, one can envisage such changes at best in an unspecific way.

For this case, however, the required generalization of our theory would have some explanatory merits. As one may easily see, cost-free and only doxastically important conceptual refinement is always not bad and usually useful; this fact is formally perfectly analogous to the usefulness of cost-free observation. This rationalizes our everlasting search for new concepts and distinctions. Of course, this result assumes that such conceptual refinements are favorable in my specific sense. However, I do not know how it could be otherwise; in contrast to doxastic changes I cannot imagine cases of purely conceptual change which should not be considered favorable. If I am right about this, the consequence is that my new optimality criterion does not come into play; the old strategic theory, though not designed for the case of conceptual change, would suffice.

The balance is different, however, in other cases of conceptual change. If the change is costly, it may not be worth the prize. If the change is a genuine one, involving also loss of old concepts (which occurs continuously, of course), the balance is wide open. And if the change has evaluative import so that the agent's utility function changes in non-trivial ways, too, then again no general verdict of the usefulness of such changes is possible.

However, in order to better understand the last assertion, we should more closely study how to deal with evaluative or motivational changes from a strategic point of view. Let us turn, hence, to the third and last group, dividing again into three parts (or four, if we include the next section).

(7) The first kind of change in the intrinsic utilities to be discussed is due to the frame-relativity of the intrinsic utility function of a decision model. The intrinsic utilities relative to the frame actually considered may really be expected utilities relative to a wider frame left implicit. So, changes of these utilities may result from all kinds of doxastic or evaluative changes within the wider frame, and hence the acceptability of these utility changes completely depends on the acceptability of the changes within the wider frame. This case of *apparent intrinsic utility changes*, as I shall call it, is therefore derivative to all the other cases here discussed.

(8) So, the interesting case is the case of *genuine intrinsic utility changes* in which the changes of the intrinsic utilities relative to the frame considered may not be understood as resulting from changes within an implicit wider frame. The first thing to be stated is that such changes are obviously pervasive; the only excuse for neglecting them is that they are usually slow and hence seldom relevant in short-term decision situations. The second thing to be stated is that, in contrast to doxastic changes, there is no general theory about such genuine intrinsic utility changes; there are just a lot of every-day experience and common-sense knowledge and a lot of more or less controlled observations and rough generalizations of motivational, developmental, and social psychology. This may be adduced as a further excuse for neglecting them; but it is none within the strategic point of view taken here, because this point of view does not require us to have an objective theory about the agent's genuine utility changes, but only to model the agent's subjective theory about her future utility changes, whatever it is, and this may be done with the fully strategic decision models defined above.

Having made these basic remarks, let me list several examples from our every-day experience. There are changes due to maturing and aging; consider, e.g., our basic drives for nutrition, sex, motion, sleep, etc. in this respect. Such changes may also be pathogenic. There are changes of taste in all areas of life due to fashions, habitualization, education, or self-education; to take a trivial example, I hated beer up to 35, only then I learnt to like it. There are all kinds of clinic addictions and phobias which develop and increase and may be cured again, and there are many changing degrees of addictive and phobic attitudes between normality and clinicality. We get bored and tired and thus lose old desires; we try out new things, because of need or curiosity, and may thus acquire new desires. We educate people not only in order to teach them knowledge, but also in order to shape their motivation and personality;

and there are many travesties of education like talking people into having certain desires, subconsciously creating specific longings, stupefaction, even brain-washing, etc.

Which of such changes are favorable and which are unfavorable? It is hard to say in view of the diversity of changes, and opinions may legitimately diverge. I content myself with three more or less general observations.

For instance, the sober attitude, presumably, towards changes due to maturing and aging is to consider them neither favorable nor unfavorable (see (8c) below for the behavioral consequences of such an assessment); but there are many who don't think so and who fight against such changes, they don't want to get old.

More importantly, there is a general tendency towards culturalization; getting more sophisticated, in one's activities, likes and dislikes, aesthetic preferences, etc., is favorable, getting more primitive is an unfavorable fall-back.³³ But there is also sophistication ad nauseam, evoking the strong desire to return to the roots. The Aristotelian order of lower and higher virtues is highly relevant here as well; getting governed by higher virtues is favorable, sinking lower is unfavorable. In the wake of this it has always been a major topic in moral philosophy whether the possible human desires can be classified in some more objective way as more or less valuable, more or less perfect, or more or less morally approvable.³⁴ Those availing of such a classification gain a natural criterion: moving to better desires is favorable, moving to worse is unfavorable.

Another point of general significance is whether the changes are recognizably forced upon the agent from outside, whether they occur naturally, or whether the agent has initiated or at least tolerated them, though she could have intervened; we have the inclination to judge changes of the first kind unfavorable, those of the second kind not unfavorable, and those of the last kind favorable. This inclination is exemplified by the fact that the advertisers take great pains to keep the advertising effect, the change of preferences, as subconscious as possible. In a broader way, the point shows itself in our attitude towards education. We adults are quite selective in allowing to get educated, and teachers fight about this with their pupils. And when suffering from ennui we still want to fill our emptiness by ourselves and are reluctant to let others do it for us. In other words, the autonomy of the evolution of

³³ McCain (1979) is one of the few references attempting to integrate this phenomenon into economic theory.

³⁴ See, for instance, the doubts raised in Mackie (1977) and the responses in Honderich (1985).

one's motives and desires is an important criterion of superiority and inferiority. Indeed, it is here where the notion of the freedom of will assumes a practically highly important sense, though it is not the one which continuously puzzles philosophers.³⁵

These observations indicate that the question which genuine utility changes are favorable and which are not touches upon deep and broad matters we cannot dwell upon here. It should be emphasized, however, that these matters are usually discussed as normative issues, which may be taken in a moral sense or as concerning individual wisdom or perfection or, maybe, as concerning rationality, though in a wider sense, then, than discussed here (this is why I avoided using "rational" in the preceding paragraph). This normative discussion is utterly important, and perhaps it succeeds in systematizing judgments of superiority and inferiority. Indeed, it is an important secondary intention of this paper to assist this discussion by providing a precise framework for developing the consequences of possible answers to the normative issues. However, one must also be aware that this discussion is only of indirect concern for us because we may take here as given which utility changes people actually judge to be favorable and which not, independently of whether or not these empirically given judgments conform to normative considerations. And certainly, these empirically given judgments can only be generalized in such a rough way as indicated above.

So, the question remains: Which recommendations does our optimality criterion give to an agent facing such genuine utility changes? Let us distinguish three cases:

(8a) I begin with the elementary case considered in the path-breaking paper of Strotz (1955/56): Suppose you firmly believe that your utility function will change in a certain way depending on the actions you have taken, and suppose all changes are unfavorable. Then the optimality criterion says that you should first determine the set of all courses of action conforming to your predictions of your future optimal behavior – these predictions being conditional on the action presently considered – and that you should then choose that course of action from this set which has maximal utility from your *present* view point. This is exactly the procedure which Strotz (1955/56) describes as *consistent planning*, and surely it is from there that I

³⁵ I say a bit more about this in my remark (6) of section 6.

have borrowed one crucial ingredient of my optimality criterion. So, from my point of view Strotz implicitly dealt with the case of purely unfavorable changes.³⁶

A slight complication of the case arises when you have some influence on whether or not an unfavorable utility change will occur.³⁷ In this case the same holds as in the case of unfavorable doxastic changes discussed in (4): you are well advised to minimize the risk of undergoing such an unfavorable change (unless this is not too costly). This is doubtlessly in accord with intuition.

(8b) Matters are simpler in the opposite elementary case where you firmly expect to undergo only favorable changes. In this case there is no need for reevaluation; as in the observational case you can rely on your future assessment already now. Hence, it is also clear what the strategic recommendation is in the more interesting case where you can choose, or at least partially control, between staying in the same or moving into a superior decision situation: the option to stay has to be reevaluated from the superior point of view; thus, the optimal substrategy after staying is not better and likely to be worse than the optimal substrategy after moving; and therefore, you should maximize the probability of the favorable change.

(8c) Perhaps the most interesting of the elementary cases is the one where you either firmly expect, or can weakly or strongly influence, to get into an incomparable decision situation. Such a change can take on various forms. One form is that it is actually a complex change composed of several favorable and unfavorable changes. In this case, the relevant superior point of view is the one which would result when only the favorable changes were to occur; this case has already received sufficient attention. The other form is that the change to the incomparable situation is not so composed; we have found this form to be plausibly exemplified by utility changes due to aging or maturing. In this case the changed situation is its own positive counterpart; thus, again, there is no need for reevaluation, you should assess what is

³⁶ It should be noted, however, that Strotz saw himself as dealing not with changing tastes, but rather with so-called time preferences, i.e. with the phenomenon that we discount future utilities from our present point of view according to some monotonically decreasing discount function and that our discount function may change in time simply because the presence shifts in time. (Such time preferences, he shows, do not give rise to temporal inconsistency if and only if the discount function is exponential.) From my point of view, however, any change in the relative weights of the utilities of the various stages of an temporally extended consumption plan is a change in the intrinsic utility function of the decision model at hand, and therefore I do not separately deal with changing time preferences as opposed to other utility changes.

³⁷ In the preceding paragraph your actions determined only which of various unfavorable utility changes is to occur. Hence, the complication already transcends the conceptual resources of Strotz' account.

optimal in that situation by its own standards. The difference to (8b) emerges only when you have some control over the change, because in contrast to (8b) the positive counterpart of the unchanged situation is the unchanged situation itself. Hence, also the positive counterparts differ in an incomparable way in this case, and therefore your choice to seek or rather avoid the change depends on the expected utilities in the two situations; you should go for the higher one.

To a good extent, my trust in the satisfactoriness of the optimality criterion I have proposed derives from the fact that the criterion gives, I think, plausible recommendations in these elementary and intuitively perspicuous cases.

(9) There is considerable economic literature on *addictions*³⁸ and apparently much less on what is, in a sense, the opposite phenomenon, namely *anxieties or phobiae*. I cannot comment on this literature in detail, though this would certainly be a useful thing to do. But I would like to point out at least that the account of strategic rationality given here seem to allow a richer and more unified picture of these phenomena. I shall restrict my remarks to addictions, but similar remarks apply, *mutatis mutandis*, to phobiae as well.

What the literature on addictions offers are mainly attempts to rationalize the development of addictive behavior. Their basic ingredient is a law of utility change with respect to a certain commodity which says how the subject's present utility for the commodity depends on the subject's past use or consumption of the commodity.³⁹ The subject may now behave in a naïve or short-sighted way with no insight into the law of utility change, thus maximizing expected utility from moment to moment, and may thus get rationally trapped into addictive behavior. Or one may assume that the subject has some insight into the law of utility change and that it follows some rule of more sophisticated choice; then a different behavior results which, however, may be still described in a way as addictive. If elaborated in detail, such models become theoretically quite intricate, and their substance is enhanced by the fact that they immediately stand to empirical confirmation and refutation insofar they have to fit observable addictive behavior.

³⁸ See Becker, Murphy (1988) and the references given there.

³⁹ The utility change may be intrinsic as in the case of genuine addictions for drugs, or it may be extrinsic and explainable by other desires; for instance, Microsoft creates an "addiction" to Windows, not because its intrinsic charm would be ever more irresistible, but because it is made an ever more indispensable part of ever more applications.

Now, the crucial point is this: Whenever it comes to model sophisticated choice, the addiction is treated as something negative, which is to be fightened and slowed down or even to get rid of. Or in my terms: the more addicted state is treated as inferior to the less addicted state. The only exception I know of is McCain (1979) who observes that there are also positive addictions⁴⁰ which behave in a different way than the negative ones.

My first remark on this is that my account is in principle capable of dealing with both cases at once. Whether a subject assesses his possible future addiction as a negative or a positive one, entirely depends on her relation of superiority/inferiority. My fully strategic decision models are sensitive to this assessment, and hence it would be interesting to check the extent to which my general decision rule agrees with the proposals made for the more specific cases.

My main remark, however, is that I find the separation into negative and positive addictions much too coarse. The attitude towards addictions seems in general to be quite ambiguous. What is bad about being a workoholic? There are not so many reasonable things to do in life promising lasting satisfaction. Why not rush then into hard work? That is how workoholism appears to be positive. But when it takes on a certain degree of persistence and exclusiveness one may as well start feeling negative about it. It makes for a reduced personality, and in extreme degrees it gives a very screwed or no safisfaction at all and makes the subject suffer. Or take alcohol: There are clear cases of alcoholism from which a small percentage of people suffer. But a much larger portion of people habitually and continuously consumes alcohol in a way which may be classified as addictive as well. Good or bad? This is not only a medical question about the long term effects of continuous doses of alcohol which are in fact heavily disputed. It is a question of how one should evaluate one's more or less strong dependence. I think there is no general answer; everybody has to find his own answer according to how pleasant or unpleasant he feels his alcoholic dependence to be.

However, if addictions have indeed this ambiguous character, the theoretical lesson is that it is not good enough to have one model for negative and one for positive addictions. Rather, one needs an account which integrates both kinds of attitudes towards addictions and even allows to consider such things like the change of a positive addiction into a negative one. And it seems that my framework at least of-

⁴⁰ His somewhat traditional examples are classical music and good wines.

fers the means for developing such an account. It is clear, however, that these general remarks are very far from actually doing this.

5. Applying the model to iterated Prisoner's Dilemma and to iterated Newcomb's Problem

I would like to suggest a further application of my model of strategic thinking involving intrinsic utility changes. It is different and more tentative than the previous ones insofar as it will assume an active rational utility change in a specific sense to be explained, and therefore I devote a separate section to it.⁴¹ The application I want to attend to primarily is the prisoner's dilemma (PD) the relevance of which to almost all areas of practical philosophy is unsurpassed; but I shall indicate how my remarks carry over to Newcomb's problem. Let me briefly resume my point of departure:

In the one-shot PD I am prepared, for the time being⁴², to accept that the only rational solution is defection; analogously, the only rational thing to do in the one-shot Newcomb problem is to take both boxes. In both cases, if you had cooperated or taken only one box, you may rightfully regret not to have chosen to get more. Thereby I am firmly on the grounds of causal as opposed to evidential decision theory. The decision maker has only an absolute, act-independent subjective probability for the other player's cooperation or the predictor's prediction, and whatever the probability is it is rational for the decision maker to take the dominating action, i.e. defecting or two-boxing. If that probability were act-dependent, that would express the decision maker's thought to have a causal influence on the other player's behavior or on the predictor's prediction – which, *ex hypothesi*, he denies and excludes.⁴³

⁴¹ This section is almost identical with sections 5 and 6 of Spohn (1999).

⁴² But see footnote 55.

⁴³ This is roughly how I argued in Spohn (1978, sect. 5.1). The basis of the argument is a probabilistic theory of causation which entails that exactly one of the four combinations of probabilistic and causal dependence and independence of the prediction (or the other player's action) on/from the decision maker's action is impossible, namely the combination of probabilistic dependence and causal independence – which is just the case which Nozick (1969) and many following him have found troublesome. Within the framework of directed acyclic graphs Meek and Glymour (1994) give an account of intervening or deciding, as opposed to prediction, which embodies the very same conclusion.

Problems start, however, with the iterated PD; indeed, I take it to be a scandal of normative rationality theory that there still does not seem to be a fully rational account in favor of cooperation. Let me briefly sketch ten different views on the iterated PD. Each view contributes a highly illuminating idea; their collection shows the tremendous intellectual challenge the iterated PD continues to present; however, even their collection is, I think, not fully satisfying.

(1) Intuitively, the case appears quite clear. If two subjects play PD many times and do not manage to set up cooperation, but are caught in defection, they are terribly silly. And it is not so that there is only collective silliness, due to a tragic conflict between individual and collective rationality. Rather, they seem to be individually accountable for their failure. At least one of them must have been individually irrational; the other one may have been irrational, too, or she may have been rational, but unable to do better than defect against her silly opponent. Intuition equally clearly tells what is rational in the iterated PD: namely to start and maintain a pattern of mutual, conditional trust and kindness which secures long and stable cooperation, relative to which disintegration in the final plays would be annoying, but negligible. This intuition is strongly confirmed by the computer tournaments of Axelrod (1984) in which the tit-for-tat strategy which instantiates this intuition in an exemplary way was most successful.

However, the attempts to back up this intuition by a theory of rationality are utterly frustrating. The central cause of all frustrations is, of course, the famous backward induction argument purporting to show that in the finitely iterated PD the only equilibrium strategy for both players is always to defect. Let us see what people have done about it.

(2) One idea is to move into the context of evolutionary game theory. Here, whole populations are occupied with playing PD, and one can choose plausible setups in which the cooperative parts of the population turn out to be much more successful than the defecting ones so that society evolves to consist mainly of cooperative individuals.⁴⁴ However, this move, illuminating as it is, simply changes the topic; we wanted to learn about individual rationality, but evolutionary game theory does not teach us in this respect.

(3) In practice, there is a simple and usually effective method to make cooperation individually rational: we legitimate an authority to reward cooperation and to

⁴⁴ Axelrod, Dion (1988) briefly present the intricacies and ramifications of the evolutionary treatment of PD.

punish defection. This is not always easy and often hardly feasible, but it considerably moderates the detriments of the PD structure. Theoretically, however, this means changing the utility functions of the players until there is no longer any PD; this is a way to avoid PD, not to solve it.

(4) So, let us look at the twin of PD, Newcomb's problem. Here, evidential decision theory seems to offer a viable rationalization of taking only one box.⁴⁵ The dominance argument is thereby invalidated and backward induction deprived of its basis, and thus one may think of carrying over this rationalization to PD. I cannot engage now into the ramified argument between evidential and causal decision theory⁴⁶; let me only express my conviction that this move is of no avail: either, evidential decision theory gets the causal relations right like in Eells (1982) and recommends two-boxing; or it neglects causal relations or gets them wrong and is therefore inadequate. Indeed, I am surprised how small the impact of the heated philosophical discussion has been in economics, in game theory, and elsewhere; evidentialism seems to be a philosophical, but only weakly contagious disease.

(5) Davis (1977) argues that the players should decide in PD according to a mirror principle saying: "If two rational agents have the same evidence and preferences, they will make the same (nonrandom) choice." (The name and the phrasing of the principle are due to Sorensen 1985, p. 158.) We may assume that both players firmly believe that this principle holds and that they satisfy its premise. Thus they are certain to do the same, and then cooperation emerges as the only rational alternative even in the one-shot case.

Is there something wrong with the mirror principle? No, I think it is important to maintain it. However, if its acceptance by the players is taken with Davis (1977) as entailing their neglect or denial of the causal independence of their actions, then we are back at the evidentialism just discarded. If it is not so taken, then it becomes clear, I think, that the mirror principle is still incomplete; it does not say anything about the rational mechanism leading the agents from given evidence and preferences to a certain choice. If that mechanism is standard game theory, both players will expect themselves to defect. Whether there is another account of rationality entailing cooperation remains the crucial question which is not answered by the mirror principle and which I want to tackle in the next section.

⁴⁵ See, for instance, Gibbard, Harper (1978).

⁴⁶ See, for instance, the papers collected in Campbell, Sowden (1985).

(6) In Spohn (1982, pp. 254-6) I argued that the standard argument for equilibrium behavior proceeding from the common knowledge of the game situation is incomplete in a very similar way. And indeed there are a lot of reasons for finding fault with the standard Nash equilibrium.⁴⁷ Might this be a way to escape the backward induction argument? No; however serious the doubts are about equilibriums in general, they seem inappropriate in the finitely iterated PD since the backward induction argument shows that always defecting is not only an equilibrium, but indeed the unique (weakly) rationalizable strategy in the sense that only it survives the iterated elimination of weakly dominated strategies. Thus, mutual knowledge of the utility functions and the Bayesian rationality of the players is essentially sufficient for establishing continued defection as the only rational option.

(7) So, perhaps the crucial fault lies in the logic of the backward induction argument itself? This doubt has been raised by Binmore (1987), Bicchieri (1989), and Pettit, Sugden (1989) and denied, for instance, by Aumann (1995). I side with Aumann, but surely one should scrutinize this discussion much more carefully than I can do here. Let me only add that as soon as one grants backward induction to hold under certain idealizations, however strong, the problem remains. If cooperation in the iterated PD can be rational for reasonable and well-informed players like us, it should be so all the more for perfectly rational and perfectly informed players; one cannot be content with allowing an exception and prescribing defection in the strongly idealized case.⁴⁸

(8) The most substantial game theoretic contributions are still to be mentioned. One line of thought is to assimilate the very often iterated PD to the infinitely iterated PD. Then one may adduce the rich battery of so-called folk theorems⁴⁹ showing that in the infinitely iterated case there are infinitely many more or less cooperative equilibria. This is an ingenious and very sophisticated observation. But it is obviously not fully satisfying, since it imputes to the players the clearly and knowably false assumption of infinite repetition.

One may, however, interpret the folk theorems in a different way: In one variant of these theorems the utilities in the future plays are discounted by some factor $\alpha < 1$, and this discount factor may also be understood as expressing the players'

⁴⁷ See, for instance, the diagrams of van Damme (1991, pp. 335f.) showing the impressive ramifications of the equilibrium concept.

⁴⁸ As does Sobel (1993, sect. 6).

⁴⁹ Cf., for instance, van Damme (1991, ch. 8) or Osborne, Rubinstein (1994, ch. 8).

subjective probability in each play that there will be a next play at all, so that the probability for an infinity of plays is in effect 0. However, I find even this interpretation implausible because there is still a positive probability for any finite number of plays. The plausible assumption would be that we all are sure to play PD at most, say, a million times in our life and very likely much less; and this assumption turns the strategic situation into a finitely iterated PD. Hence, cooperation should be rationally possible also in this case.

(9) In conversation, Teddy Seidenfeld proposed to me another variant which drastically changes the picture: make the continuation of the game in some way dependent on past cooperation; there may or may not be an upper limit to the number of plays. This idea has already been fruitfully applied by Feldman and Thomas (1987) in the context of evolutionary game theory. Its point in the context of individual decision making is obvious: this variant set-up provides for a simple and theoretically sound rationalization of cooperation in standard decision theoretic terms and avoids the devastating backward induction. This is a beautiful idea, but it provides only a partial solution and dissolves neither the original problem nor the desire to solve it as well.

(10) The idealizations required for backward induction, i.e. the relevant assumptions of mutual knowledge, may well fail, of course. This is the entry of the perhaps most interesting attempts to establish a cooperative solution. The catch notions are the (trembling hand) perfect equilibria of Selten (1975) and the sequential equilibria of Kreps, Wilson (1982). Very roughly, the idea pursued here is that a rational strategy must define rational behavior even for situations which can only be reached if some players behave irrationally and that one should always expect with a small probability that irrational behavior occurs intentionally or unintentionally. Taking these things into account may further cooperation in the following way: I cooperate in the first play, because my hand trembles or maybe because I follow a sophisticated plan. My partner is surprised, but then starts thinking how my perceived irrationality can be explained and maintained, and perhaps he reaches the conclusion that he should cooperate as well in the second play, and so on; cooperation may thus be the perfectly rational continuation of a somehow irregular or irrational beginning.⁵⁰

⁵⁰ A precise story is told by Kreps et al. (1982).

This picture may be realistic, but from the point of view of normative rationality theory it seems distorted. The normative intuition which demands compliance is that cooperation is rationally possible and indeed rational, and that it must be so without any help from direct or indirect gaps or deficiencies in rationality; it seems not good enough to show how cooperation can emerge as a form of bounded rationality.

So the suggestion in particular from (7) and (10) is that it is the normative rationality theory itself which needs getting reformed, and my brief summary was, I hope, fair in suggesting that no working idea for this reform seems available.

However, in section 3 I have already proposed a reform of rationality theory. Does it help, perhaps, to illuminate the present problem as well? Yes, I think it does in a certain way – a way which may seem cheap or miraculous; but it would be surprising, on the other hand, if the solution would have to be very complicated or sophisticated. So, here is the line of thought I want to propose:

Intuitively, we would reason as follows in the finitely iterated PD: Considered in itself I should defect in the first play because I have no influence whatsoever on what you do in the first play. However, when I cooperate now this may have a positive impact on, i.e. raise the probability of your cooperation in later plays. And that is how cooperation in the first play may have maximal expected utility. So far, so good. But how could I raise the probability of your cooperation in the second play? By the same consideration, your cooperation in the second play can only get a positive probability if you think that it raises the probability of my cooperation in the third play. And so the hope of raising the probability of cooperation is deferred to later and later plays until the final play where we know already that it will be badly disappointed. Hence, there is no rational hope in making cooperation likelier, and thus the intuitive reasoning fails.⁵¹ Of course, this is again nothing but a form of the old backward induction argument.

Let me put this impossibility in a somewhat different way: It is constitutive of PD that I do not believe in a correlation between our choices in the first play; your choice is just an independent, i.e. causally and (hence⁵²) probabilistically independent state of the world for me. But somehow we would like to believe in a probabilistic correlation between our actions in later plays. How could we have this belief?

⁵¹ By contrast, however, this reasoning would be perfect in the scenario considered in point (9) of this section.

⁵² See footnote 43 above.

The first difficulty is that as a decision maker pondering about my future actions I do not have any subjective probabilities for these future actions of mine; I determine the best or rational course of action, and then we may or may not add the epiphenomenal belief in that course of action. Indeed, I still think that this is an important point which Savage (1954) got right and Jeffrey (1965) wrong. From the point of view of rationality theory it is only the above generalized strategic thinking which allows us to have a probabilistic assessment of the actions considered, namely by assuming subjective probabilities for getting into various future decision situations which rationality theory must view as complete deterministic causes for the actions taken in them.

Hence, we can believe in a correlation between our actions in the later plays of PD only if we believe our decision situations in the later plays to be correlated. So, I have to imagine myself being in various possible decision situations later on, say, in the second play. How could they vary? Apparently only by containing varying beliefs, i.e. varying subjective probability functions. But how then could the possible decision situations in the second play also vary in their optimal actions, as they must when a correlation between these actions is to emerge as well? Only by containing varying assumptions about the correlations in the third or in later plays. They cannot contain such assumptions concerning the second play itself, because then I can no longer believe our choices in the second play to be correlated; I have to believe then into the causal and (hence) probabilistic independence of these choices. In this way a present belief in future correlation can only derive from a future belief in still more distant correlation, and again the castle in the air collapses in the last play. Thus, if belief in correlated future action presupposes belief in correlated future decision situations, and if these future situations differ only in their subjective probabilities, the backward induction argument strikes again, and there is no rational way to entertain such a belief.

However, the last conclusion suggests to consider a further possibility: namely that the future decision situations we would like to believe to be correlated differ also in their utility functions. How might this come about? This possibility seems to violate the very set-up of the iterated PD. But no, there is, I think, a way of rationalizing this suggestion.

Let us start from the supposition that I believe that our actions are correlated, not in the first play, but in the second and later plays. As a consequence, I realize that we are caught in a continued perspectival trap by the iterated PD which consists

in the fact that the actions in later plays which I now believe to be correlated cannot seem correlated to me at the later time of choice; at that later time I can view your action only as a causally and hence probabilistically independent state of the world.

Now, the theory of section 3 enters the argument. For, if I perceive these later plays as a trap, I judge these later decision situations as we have conceived them so far to be inferior in the specific sense discussed earlier. The relevant superior situation from which to assess this inferior situation is the one in which my cooperation and my defection in the later plays are reevaluated by receiving an additional utility or disutility which makes their overall utility correspond to the expected utility they would have under the presently assumed correlation with your actions in the later plays.

It is not implausible, I think, to apply here the superiority/inferiority distinction in this way. On reflection, such a trappy structure turns out to be not unusual; quite often the formation of our decision situations is in some way negatively or counterproductively correlated with the external circumstances. For instance, a strong desire to eat often arises due to some more or less subconscious frustration which, unlike hunger, does not vanish by eating; hence, a situation in which such a desire for food is present is inferior to the same situation without this desire. Similarly, there is a not uncommon tendency of men to fall in love, i.e. to assign a high utility to being related with women who are completely reserved and to ignore women who are obliging. Sometimes, men have this tendency because they do not really want to get close to women. But others severely suffer from this tendency; and for them the situation of being in love with an unapproachable woman and being disinterested in a responsive one is inferior to the reverse situation. These examples differ, however, from iterated PD, and iterated Newcomb as well, because they embody a trappy desire formation, whereas it is belief formation which is trappy in iterated PD and Newcomb. Indeed, it is so in a perfectly schematic way simply due to my moving in time; at any time the other player's next choice, or the predictor's next prediction, is probabilistically independent from my next choice only from the present, but not from any earlier point of view. This remarkable feature should, however, not distract from the similarity to the other cases.

So, let us return to PD. What is the point of thus applying the superiority/inferiority distinction? It does not yet seem to help: As I have explained in the first part, the relevant superior situation is so far only a hypothetical situation, and if the decision situations I actually expect to reach are those in the trap, the only consistent

strategy is still always defecting; hypothetical reevaluation alone cannot change this. So a further step is required. It says that if my superiority assessment is as explained, I should actually move into the superior situation. This may seem strange, but it is the core of the solution I want to propose. In the examples discussed in section 4 I was moved into new decision situations by external forces like observation, alcohol, and other things in a way which was not under my immediate control. In the present case, by contrast, I want to suggest that it is the pure insight into the trap-like structure of the whole set-up which should rationally move me into the superior situation with its adjusted utility function. What I thus propose, in effect, is a *law of rational utility change* which is not a mere change of expected utilities due to changes of subjective probabilities.⁵³ This resembles the usual practical solution of PD mentioned in point (3) of this section, which consists in changing the utility functions of the players from outside; my suggestion is that internal rationality alone should have the same effect as external punishment.

Now my argument may be brought to an end. The full theory of rationality is now the one amended by these laws and mechanisms. Hence, if I think that you are rational I think that you conform to this amended theory as well. Moreover, I do not firmly believe into a fixed correlation between our actions in the later plays. Correspondingly, there is no fixed superior decision situation into which I should move. Rather, how strong a correlation I assume depends on our present actions which may or may not intensify the assumed correlation. In this way, cooperation may or may not become the rational thing to do in the later plays, and it may do so for me and for you in a correlated way.⁵⁴ Rationally, however, you and I should start with believing in a strong correlation which will then be confirmed so that long-standing cooperation will indeed emerge. The crucial point of this reasoning is that thereby the present belief in the correlation in later plays does not derive from the later belief in still more distant correlation. In this way, the reasoning breaks the force of the backward induction argument.

To summarize, my argument is that the subjective assumption that the chosen actions will be correlated puts in force an enriched theory of rationality employing superiority assessments and a new law for changing utilities, and this enriched the-

⁵³ For further comments on this law of rational utility change see remark (4) in section 6.

⁵⁴ It would be interesting at this point to study the relation to the theory of correlated equilibria initiated by Aumann (1974). Indeed, my account may help explaining how a correlated equilibrium may establish without explicit agreement or external arbiter.

ory in turn makes the assumption of correlation rationally entertainable. By this kind of bootstrapping, correlation is rationally believable and cooperation thus rationally possible in a full sense, even in the finitely iterated PD.

Mutatis mutandis, these considerations should apply to the iterated Newcomb problem. This then would be my offer as a causal decision theorist to the evidentialist whose intuition I share that it cannot be rational to stay poor if one has the chance to get rich. It is not true that "the reason why we", the causal decision theorists, "are not rich is that the riches were reserved for the irrational" (Lewis 1981, p. 377); the reason is that we were caught in too narrow a notion of rationality. The truly rational does not pity himself because only the allegedly irrational are consistently (pre-)rewarded; he should be able to adapt, and my proposal shows a way how to rationally do so.⁵⁵

6. A list of further issues

So far, my exposition may have appeared to have an orderly shape, though in the previous section I was already content with giving only a suggestive sketch. In fact, however, the subject matter is haunted by numerous unclarities and problems, and I have only very insufficient conceptions of how to come to grips with them. So let me use this final section for eight miscellaneous remarks which show the richness of topics to which my account of strategic thinking relates. It may well be that the orderly shape I have preserved so far will thereby collapse completely and that only confusion remains. This, however, would not be an unusual result in philosophy.

(1) I have not made explicit so far, though it may have been clear all along, that my account thoroughly presupposes an *intrapersonal utility comparison*. The utility functions of all the possible decision situations considered in a fully strategic decision must be on one and the same scale; otherwise, my account would not make any sense at all. To be a bit more precise: the consistency constraint does not require

⁵⁵ It has been suggested to me several times in conversation that my reasoning might be applied to the one-shot case as well. Perhaps, though the causal independence of your action (in PD) or the predictor's prediction (in Newcomb) from my action looks unshakeable. It is obvious that such suggestions can only be checked after my proposal in this section is carried out with formal rigor – something which urgently needs to be done.

any utility comparison across the situations considered to be actually reached, because for satisfying it one only has to know what is optimal in each of these situations by itself. The principle of reevaluation, however, is more demanding. It relies on the expectation of the reevaluation of the considered substrategies in the relevant positive counterparts of the actual situations, and therefore the utility function of all these positive counterparts must be fully comparable, i.e. unit and origin comparable.

I am not worried about this presupposition. But it should be explicit, since interpersonal utility comparison is deemed problematic by many and since intrapersonal comparison may therefore seem equally problematic. However, I shall not go to defend my presupposition.

(2) In Spohn (1977) I argued that Jeffrey's (1965) logic of decision is flawed by assuming the agent to have *subjective probabilities for her own future actions* and proposed to return to Fishburn's (1964) modelling of the agent's subjective probabilities which I have introduced in section 2. Jeffrey (1977) opposed my criticism. Since then the issue stayed between my intuition that there is something very odd about probabilistically assessing one's own future actions and Jeffrey's intuition that we can predict our future behavior just as well or badly as any other future events.

I granted from the beginning that I see no problem in having beliefs about one's future behavior insofar as it is not an object of one's present practical deliberation. I wanted to exclude only those actions which figure in one's decision situation as things to decide upon. Now I can do a little bit more for reconciliation. Strategic thinking allows one, of course, also to have probabilities for the actions presently under consideration. They do so, however, only under the reflective perspective of strategic thinking, only by having probabilities for future decision situations one might get into and by knowing which actions would be optimal in them. This does still not yield a probability for the first action which is to be taken now and which does not figure any more as action in future decision situations. Here, one could only trivially add, as an epiphenomenal consequence of one's practical deliberation, that the optimal option has probability 1, leaving 0 for all others.

This is how probabilities for one's own future actions may enter the picture. Still, all this is only partial reconciliation. The basic difference remains, I think, insofar as I insist that probabilities for one's own actions emerge only as a result of

the practical deliberation by self-applying rationality theory and should not enter into the deliberational model from the beginning.

(3) Though it would be desirable I cannot accomplish here an extensive comparison between my proposal and the many ideas in the literature which are directed to the problems discussed here. Let me confine myself to two brief remarks concerning sophisticated and resolute choice.

If there is something like a standard account in the field, it is certainly the theory of *sophisticated choice* which is basically characterized by what I called above the consistency constraint.⁵⁶ This constraint is a severe one, and there are many situations in which it already suffices to determine an optimal strategy. However, there are many situations – I have discussed many of them in section 4 – in which it does not suffice; and then it is important to close the remaining gap. Here, my principle of reevaluation provides an elaborate proposal, and insofar my account may be seen as a refinement of the theory of sophisticated choice.

However, this is not the only perspective one may take on it. McClellenn (1990, sect. 11.3) has forcefully argued that sophisticated choice may lead to intuitively unreasonable results and that what he calls *resolute choice* seems often to do better. Here, resolute choice consists in behaving naïvely, as it were, for a certain stretch of time, i.e. in deciding in advance for a whole course of action extending through that stretch and implementing that course without looking left or right, i.e. without reconsidering in between. Resolute choice looks reasonable just in cases where sophisticated choice at each decision node drives one to an end point which is suboptimal, indeed strictly dominated among all end points there are.

A nice example is the so-called toxin puzzle invented by Kavka (1983), in which I am generously rewarded for having the firm intention at midnight to drink a glass of toxin on the next day which makes me painfully sick for several hours, but has no lasting effects whatsoever. The point is that I am rewarded for having the intention, afterwards I can just forget my intention and abstain from the toxin. But then it seems I cannot even credibly form the intention and thus have no chance of getting the reward which would by far outweigh the pain. This is a case for a reso-

⁵⁶ See Strotz (1955/56), Pollak (1968), Peleg, Yaari (1973), Hammond (1976, 1988), Yaari (1977), and Schelling (1984). These references and in particular the latter book would deserve more careful comparative remarks.

lute choice which recommends forming the intention, sticking to it, and finally actually drinking the toxin.

It seems to me that there are two ways for my model to encompass resolute choice. The first is only a formal trick. It consists in observing that it is not decided in advance where the choice points in a strategic decision tree really are. My model supposed that there is always one action taken at on step of time. But it did not say how large such a step is, whether it is not rather a stretch of time, and whether the action to be decided is not rather a whole course of action extending through that stretch. Hence I can always conceive the decision tree in such a way that resolute choice reduces to ordinary choice.

This trick may, however, appear illegitimate because one needs to look at the more fine-grained decision tree before seeing with which coarsening one might be better off. Therefore I have a second suggestion, though only very tentatively. Critics have found resolute choice incredulous because there is no way of committing oneself to marching through to the optimal end point; one is supposed to do so, so to speak, by a sheer act of will. But I think the case is quite similar to what I have said about the iterated prisoner's dilemma. There as well as here, in the toxin puzzle, for instance, the choice situation is trappy; the decision rules lead the subject into behavior which is suboptimal by her own lights. Thus I have proposed that insight in that trappy structure should rationally move the subject to change her utility function in an appropriate way so that the reasonable behavior ensues. Certainly, critics will not be satisfied; such utility changes are as miraculous as is resolute choice. But my suggestion may show a way of incorporating and thus of rationalizing resolute choice within my account. At the same time, it shows that it is perhaps too narrow to classify my account simply as an elaboration of sophisticated choice.

(4) The previous paragraph showed once more the important role the envisaged *law of rational utility change* plays in my account. If it really does so, then one should expect, however, that it has a much wider application beyond the special PD set-up. This is indeed the case, I believe. It seems to be generally true that, when I am apparently stuck in an inferior position, I should rationally change my evaluations so as to actually reach the so far only hypothetical superior point of view. This was my suggestion in the iterated PD, in the iterated Newcomb, and now also in the toxin puzzle. But it applies to addictions (and phobiae) as well. If I realize I am (getting) addicted I should assign an evaluative malus to future addictive behavior so

that total abstention or modest consumption has the higher utility. Or look again at the two examples introduced in the previous section. If I realize that I am stuffing myself out of frustration and that my desire to eat thus puts me into an inferior situation, eating for that reason should rationally receive a negative evaluation. And if I realize my self-frustrating longing for unapproachable women, this should dampen my longing in future cases. And so on.

In the latter cases the utility change may seem even more mysterious as in the former ones. How should I be able to change my utilities effortlessly, as it were, by pure insight? This sounds less mysterious if we observe a thorough-going ambiguity in the notion of utility. As Kusser (1989) has made convincingly clear, we have to distinguish evaluation, motivation, and satisfaction, as she calls it. Our evaluations of certain states of affairs represent our conscious desires or utilities as they are reflected in sincere speech; the motivating force of states of affairs manifests itself in our actions; and the potential of satisfaction of states of affairs shows up in the pleasantness and unpleasantness of the resulting states of the agent. Ideally, and perhaps also usually, these three aspects are congruous so that one may neglect these distinctions (as the received decision theory did). However, the three aspects may, and often do, diverge, and these divergences give rise to various interesting, practically relevant, and theoretically consequential phenomena.⁵⁷

This observation is relevant for us insofar as the law of rational utility change I have proposed makes sense only when it is construed as referring to the evaluative aspect of utility. So construed, utility change seems easy; my judgment that my evaluation should be different is tantamount to a change of my evaluation. But there is no presumption that my motivation automatically follows my evaluation. In the case of PD and Newcomb motivation should do so without any difficulty at all; in the toxin case it is less easy, since it really costs an effort to drink the toxin; and in the case of addictions or eating out of frustration the insight into the inferiority of the situation and the evaluational change alone may be more or less powerless to change one's motivation.

Once this presumption is cleared away, my law of rational utility change may look much more plausible. On the other hand, this remark sets the task to take precise account of the different aspects of utility within the theory of strategic rationality, a task I have not touched at all in this paper.

⁵⁷ Cf. Kusser (1989, pp. 108ff and pp. 164ff.) and Kusser, Spohn (1992).

(5) There is another apparently remote connection I would like to point out. Frankfurt (1971) has initiated a fascinating philosophical discussion about the importance of *second order desires or preferences*; he argued that second order desires are important for free will and for personhood. I am not prepared to contribute anything to this discussion. The only comment I want to make (but see also remark (7)) is this:

Due to its nature and its setting, the discussion has hardly received any formal clarification. There is hence no account of how precisely second order desires interact with first order desires and how they eventually influence or direct behavior. But such an account is certainly missing; without it this discussion hangs in the air. Now it seems to me that my model of strategic thinking may be used as a proposal in this direction. A subject's superiority/inferiority relation can clearly be conceived as capturing her second order preferences, and thus my model determines specific behavioral consequences of these second order preferences within the strategic context. It would certainly be useful to work out this connection in a more substantial way. That it exists may, however, already have been apparent in section 4, (8), when I listed a number of points determining our superiority relation insofar intrinsic utility changes are involved and which are clearly in the vicinity of this discussion.

(6) *Freedom of the will* is too large a topic to be dealt with now in any responsible way. But I would like to point at least that this paper is implicitly closely related to this topic insofar as it opens three different levels on which freedom of the will may be located.

The basic level is related to the issue (2) of this section., i.e. to the principle that any model of practical deliberation should not contain subjective probabilities of the agent for her own future actions. Thus, the agent's future actions are not objects of her beliefs. But, of course, she has beliefs about other things conditional on her actions. This entails that she sees her actions as possible causes of other events, but themselves as *uncaused*; her practical deliberation does not contain or refer to any causes of her actions. This is one and, as it seems to me, the basic sense in which these actions are called to be free.⁵⁸ This remains true even in the strategic perspective propounded here, because in this perspective at least the next action to be

⁵⁸ I have put forward this idea already in Spohn (1978, p. 193).

chosen is not an object of the agent's beliefs and thus seen as uncaused and as free in the relevant sense.

On the second level we leave the subjective perspective of practical deliberation and take an external point of view. Here we can state the agent's *freedom of action* in the sense that she can do what she wants, i.e. what seems best to her in the given situation. Freedom of action in this sense obtains trivially; everybody, even the man in chains, has it as long as he acts at all. This sense becomes more substantial only if we start measuring somehow the size of the agent's set of options which is small for the man in chains and large for the tyrant who is subjected only to physical and not to civil laws. Thus we have degrees of freedom on that level and a serious ethical discussion about the allocation of these degrees of freedom and thus, for instance, about fundamental human rights.

On the third level we stick to the external perspective and consider the agent's *freedom of her will*, that is within the decision theoretic picture, the freedom of her desires or her (intrinsic) utilities. This freedom cannot consist in the ability to choose one's desires as one wants, i.e. as they seem best to one; one cannot choose one's desires or utilities as one can choose one's actions. I am not very clear about what else freedom of the will in the third sense consists in. But one crucial aspect of it is certainly the extent to which the evolution of one's desires or utilities is under the control of one's second order desires or superiority relation. This is vague, but it should include the absence of influences on one's utilities which one cannot avoid or switch off, and it may include such things as obedience to the law of rational utility change envisaged above or the power to choose one's education.

I should not indulge further into speculation, however. It is too obvious that I have moved onto slippery ground with my remarks (4) - (6) and that a lot of clarification work is still required for knowing how to move well on these grounds.

(7) This brings me to some remarks arguing against *possible generalizations* of my account.

The first possibility follows up to my remark (5). Obviously I have treated first and second order in a different way; first order desires were specified cardinally by utility functions, whereas second order desires were represented ordinally by the superiority relation. So, why not introduce cardinal utilities at the second level, too? Would that not look more unified?

However, I do not see any prospects for this idea; in any case it does not fit into the account presented here. The basic reason is that the superiority relation functions in my account in a lexicographic way. There is no sense in compromising or weighing between an inferior and a superior situation in reevaluating which must always be done from the most superior situation available; only evaluationally equivalent situations can be weighed and mixed. This observation, if it is correct, entails not only that ordering is sufficient, but also that there is no sense in cardinally measuring the superiority relation; there is no place in the theory for superiority degrees to show up and, reversely, to get inferred.

This leads me to consider the second possible generalization. In section 2, when defining observational decision situations, I restricted the utility function U at time t_0 to the set V of action and (old) occurrence variables; only the subjective probabilities P at t_0 were to cover also the variable Δ collecting the possible decision situations at t_1 . And I stucked to this restriction in all later definitions. But why not extend U to Δ as well? Why not assume a utility for getting into a decision situation δ beyond the utilities in δ themselves? It may even appear that this might have enormously simplified my account. Could we then not replace the complicated reevaluation business by simply giving the superior (inferior) situations a sufficient utility bonus (malus)? This might perhaps also been the right way to conceive of second order desires?

However, this replacement would not work in the appropriate way. It would introduce a trade-off between superior and inferior decision situations which I wanted to avoid all the time. If I could choose, why not be a lucky idiot rather than an ever doubtful intellectual? If the superior situation of being intellectual is given only a utility bonus, it may be compromised; if the possible idiot I could be would be *very* lucky, I might start to prefer being him. This seems wrong to me. Moreover, the superior points of view then enter the practical deliberation only as weighted with the subjective probabilities for their being actually reached. Again, this seems wrong to me; as I have explained in section 3, the superior decision situations do their work in reevaluation even if the subject takes them to be only hypothetical and has no presumption to actually reach them.

But even if that simplification is misconceived, the question remains why one should not allow U to extend to Δ within my account. I have refrained from this because my account would have contained then two different kinds of second order preferences; that would have appeared at least weird. There is certainly good intuitive

sense in finding decision situations more or less valuable. But the two senses I can make of this are already taken care of in my account. In one sense, I prefer to be in a decision situation δ rather than in δ' simply because, given intrapersonal utility comparison, my utilities in δ are, on the whole, higher than those in δ' ; so, this kind of evaluation derives from the first order utilities in the possible decision situations. In the other sense, I have a second order superiority evaluation of the possible situations. But I do not see a third sense in which decision situations might be more or less valuable, and hence I do not see the necessity to liberalize my account in this respect.

(8) However, there is finally a *generalization* which seems required: Has there been an assumption of *perfect recall* built in into my account? No, in some sense even memory loss has been built in. Let me explain:

Generally, perfect recall is the assumption that the subject knows at each decision node what she has done and what has occurred to her so far. In usual game and decision trees this comes down to assuming that the subject always knows his past actions. In this specific sense, perfect recall is not part of my account; the agent may assume about her future development whatever she wants, even that she will forget what she has done.

The case is different, however, with the hypothetical superior decision situations. In them, the subject should always recall her past actions, since, as I have argued in section 4, (4), any kind of forgetting is an unfavorable change. So, here the assumption of perfect recall is entirely appropriate.

There are other things, though, than the own actions which the subject may or may not remember; and in one respect I have indeed assumed total memory loss. The reason is that in my fully strategic decision models I have assumed the agent to have a lot of beliefs about her future decision situations, but none about the past ones; the situations she has been in simply do not figure in this model. So one might think how to take into account also these kinds of beliefs.

It seems that this is not just a fancy possibility, but indeed intuitively required. Probably everyone has been in the situation that she only partially remembers the situation she has been in; say, she (partially) remembers the strategy she has started to follow, but cannot fully recall or reconstruct the reasons why; or she remembers many, but not all parts of her past decision situation and cannot recall or reconstruct the strategy that has been optimal. What to do then? Simply forget about the

remaining recollections about one's past situation and try to do as well as one can from the present point of view? Would that be reasonable even if you suspect that your past situation which you recall only insufficiently is superior as compared with your present one? (Imagine you are drunk and inclined to do a because this seems to serve your present interests best, but faintly recall that for some reason or other it was b that you had intended to do.)⁵⁹

I do not know at all how to handle such cases. It may very well be that they call for a non-trivial extension of my present account, indeed perhaps a change insofar as accommodating the new cases may allow or suggest a different treatment of the cases covered here. This is one of the many serious questions I must leave open.

Bibliography

- Aumann, R. (1974), "Subjectivity and Correlation in Randomized Strategies", *Journal of Mathematical Economics* 1, 67-96.
- Aumann, R. (1995), "Backward Induction and Common Knowledge of Rationality", *Games and Economic Behavior* 8, 6-19.
- Aumann, R., S. Hart, M. Perry (1997), "The Absent-Minded Driver", *Games and Economic Behavior* 20, 102-116.
- Axelrod, R. (1984), *The Evolution of Cooperation*, New York: Basic Books.
- Axelrod, R., and D. Dion (1988), "The Further Evolution of Cooperation", *Science* 242, 1385-1390.
- Becker, G.S., K.M. Murphy, (1988), "A Theory of Rational Addiction", *Journal of Political Economics* 96, 675-700.
- Bicchieri, C. (1989), "Self-Refuting Theories of Strategic Interaction: A Paradox of Common Knowledge", *Erkenntnis* 30, 69-85.
- Binmore, K. (1987), "Modeling Rational Players: Part I", *Economics and Philosophy* 3, 179-214.
- Campbell, R., L. Sowden (eds.) (1985), *Paradoxes of Rationality and Cooperation*, Vancouver: University of British Columbia Press.
- Csiszar, I. (1975), "I-Divergence Geometry of Probability Distributions and Minimization Problems", *Annals of Probability* 3, 146-158.
- Davis, L. (1977), "Prisoners, Paradox, and Rationality", *American Philosophical Quarterly* 114, 319-327.
- Eells, E. (1982), *Rational Decision and Causality*, Cambridge: Cambridge University Press.
- Elster, J. (1979), *Ulysses and the Sirens*, Cambridge: Cambridge University Press.
- Feldman, M.W., E.A.C. Thomas (1987), "Behavior-Dependent Contexts for Repeated Plays of the Prisoner's Dilemma II: Dynamical Aspects of the Evolution of Cooperation", *Journal of Theoretical Biology* 128, 297-315.

⁵⁹ A nice problem of this sort, the absent-minded driver, has been invented by Piccione, Rubinstein (1997) and further discussed by Aumann, Hart, Perry (1997).

- Fishburn, P.C. (1964), *Decision and Value Theory*, New York: Wiley.
- Frankfurt, H. (1971), "Freedom of the Will and the Concept of a Person", *Journal of Philosophy* 68, 5-20.
- Gibbard, A., W.L. Harper (1978), "Counterfactuals and Two Kinds of Expected Utility", in: C.A. Hooker, J.J. Leach, E.F. McClennen (eds.), *Foundations and Applications of Decision Theory*, vol. 1, Dordrecht: Reidel, pp. 125-162.
- Goldstein, M. (1983), "The Prevision of a Prevision", *Journal of the American Statistical Association* 78, 817-819.
- Hammond, P. (1976), "Changing Tastes and Coherent Dynamic Choice", *Review of Economic Studies* 43, 159-173.
- Hammond, P. (1988), "Consequentialist Foundations for Expected Utility", *Theory and Decision* 25, 25-78.
- Hild, M. (1998a), "Auto-Epistemology and Updating", *Journal of Philosophical Logic*, to appear.
- Hild, M. (1998b), "The Coherence Argument Against Conditionalization", *Synthese*, to appear.
- Honderich, T. (ed.) (1985), *Morality and Objectivity*, London, Routledge & Kegan Paul.
- Hunter, D. (1991), "Maximum Entropy Updating and Conditionalization", in: W. Spohn, B.C. van Fraassen, B. Skyrms (eds.), *Existence and Explanation*, Dordrecht: Kluwer, pp. 45-57.
- Jeffrey, R.C. (1965), *The Logic of Decision*, Chicago: Chicago University Press, 2nd ed. 1983.
- Jeffrey, R.C. (1977), "A Note on the Kinematics of Preference", *Erkenntnis* 11, 135-141.
- Jeffrey, R.C. (1988), "Conditioning, Kinematics, and Exchangeability", in: B. Skyrms, W.L. Harper (eds.), *Causation, Chance, and Credence*, vol. 1, Dordrecht: Kluwer, pp. 221-255.
- Kavka, G. (1983), "The Toxin Puzzle", *Analysis* 43, 33-36.
- Kreps, D.M., P. Milgrom, P., J. Roberts, R. Wilson (1982), "Rational Cooperation in the Finitely Repeated Prisoners' Dilemma", *Journal of Economic Theory* 27, 245-272.
- Kreps, D.M., R. Wilson (1982), "Sequential Equilibria", *Econometrica* 50, 863-894.
- Kusser, A. (1989), *Dimensionen der Kritik von Wünschen*, Frankfurt a.M., Athenäum.
- Kusser, A., W. Spohn (1992), "The Utility of Pleasure is a Pain for Decision Theory", *Journal of Philosophy* 89, 10-29.
- Lewis, D. (1979), "Attitudes *De Dicto* and *De Se*" *Philosophical Review* 88, 513-543; supplemented with postscripts in: D. Lewis, *Philosophical Papers*, Vol. I. Oxford: Oxford University Press, 1983, 133-156.
- Lewis, D. (1981), "Why Ain'cha Rich?", *Noûs* 15, 377-380.
- Mackie, J.L. (1977), *Ethics. Inventing Right and Wrong*, Harmondsworth, Pelican.
- Maher, P. (1992), "Diachronic Rationality", *Philosophy of Science* 59, 120-141.
- Maher, P. (1993), *Betting on Theories*, Cambridge: Cambridge University Press.
- McCain, R.A. (1979), "Reflections on the Cultivation of Taste", *Journal of Cultural Economics* 3, 30-52.
- McClennen, E.F. (1990), *Rationality and Dynamic Choice*, Cambridge: Cambridge University Press.
- Meek, C., C. Glymour (1994), "Conditioning and Intervening", *British Journal for the Philosophy of Science* 45, 1001-1021.
- Nozick, R. (1969), "Newcomb's Problem and Two Principles of Choice", in: N. Rescher et al. (eds.), *Essays in Honor of Carl G. Hempel*, Dordrecht: Reidel, pp. 114-146.
- Osborne, M.J., A. Rubinstein (1994), *A Course in Game Theory*, Cambridge, Mass.: MIT Press.

- Peleg, B., M.E. Yaari (1973), "On the Existence of a Consistent Course of Action When Tastes are Changing", *Review of Economic Studies* 40, 391-401.
- Pettit, P., R. Sugden (1989), "The Backward Induction Paradox", *Journal of Philosophy* 86, 169-182.
- Piccione, M. A. Rubinstein (1997), "On the Interpretation of Decision Problems with Imperfect Recall", *Games and Economic Behavior* 20, 3-24.
- Pollak, R.A. (1968), "Consistent Planning", *Review of Economic Studies* 35, 201-208.
- Savage, L.J. (1954), *The Foundations of Statistics*, New York: Wiley, 2nd ed.: Dover 1972.
- Schelling, T.C. (1984), *Choice and Consequence*, Cambridge, Mass.: Harvard University Press.
- Selten, R. (1975), "Reexamination of the Perfectness Concept for Equilibrium Points in Extensive Games", *International Journal of Game Theory* 4, 25-55.
- Skyrms, B. (1987), "Updating, Supposing, and Maxent", *Theory and Decision* 22, 225-246.
- Skyrms, B. (1990), *The Dynamics of Rational Deliberation*, Cambridge, Mass.: Harvard University Press.
- Sobel, J.H. (1993), "Backward-Induction Arguments: A Paradox Regained", *Philosophy of Science* 60, 114-133.
- Sorensen, R.A. (1985), "The Iterated Versions of Newcomb's Problem and the Prisoner's Dilemma", *Synthese* 63, 157-166.
- Spohn, W. (1977), "Where Luce and Krantz Do Really Generalize Savage's Decision Model", *Erkenntnis* 11, 113-134.
- Spohn, W. (1978), *Grundlagen der Entscheidungstheorie*, Kronberg/Ts.: Scriptor.
- Spohn, W. (1982), "How to Make Sense of Game Theory", in: W. Stegmüller, W. Balzer, W. Spohn (eds.), *Philosophy of Economics*, Berlin: Springer, pp. 239-270.
- Spohn, W. (1997), "Über die Gegenstände des Glaubens", in: G. Meggle (ed.), *Analyomen 2. Proceedings of the 2nd Conference "Perspectives in Analytical Philosophy". Vol. I: Logic, Epistemology, Philosophy of Science*, Berlin: de Gruyter, pp. 291-321.
- Spohn, W. (1998) "The Intentional versus the Propositional Conception of the Objects of Belief", in: L. Villegas, M. Rivas Monroy, C. Martinez (eds.), *Proceedings of the Congress on Truth, Logic, and Representation of the World in Santiago de Compostela 1996*, Santiago de Compostela, pp. 266-286.
- Spohn, W. (1999), "A Rationalization of Cooperation in the Iterated Prisoner's Dilemma", in: J. Nida-Rümelin, W. Spohn (eds.), *Practical Rationality, Rules, and Structure*, Dordrecht: Kluwer, im Erscheinen.
- Strotz, R.H. (1955/56), "Myopia and Inconsistency in Dynamic Utility Maximization", *Review of Economic Studies* 23, 165-180.
- van Damme, E. (1991), *Stability and Perfection of Nash Equilibria*, Berlin: Springer, 2nd ed.
- van Fraassen, B.C. (1984), "Belief and the Will", *Journal of Philosophy* 81, 235-256.
- van Fraassen, B.C. (1995), "Belief and the Problem of Ulysses and the Sirens", *Philosophical Studies* 77, 7-37.
- Wagner, C.G. (1992), "Generalized Probability Kinematics", *Erkenntnis* 36, 245-257.
- Weirich, P. (1983), "Conditional Probabilities and Probabilities Given Knowledge of a Condition", *Philosophy of Science* 50, 82-95.
- Yaari, M.E. (1977), "Endogeneous Changes in Tastes: A Philosophical Discussion", *Erkenntnis* 11, 157-196.