Decisions in Dynamic Settings[1]

Paul Weirich

University of Rochester

The expected utility of an option for a decision maker is defined with respect to probability and utility functions that represent the decision maker's beliefs and desires. Therefore, as the decision maker's beliefs and desires change, the expected utility of an opinion may change. Some options are such that their realizations change beliefs and desires in ways that change the expected utilities of the options. If a decision is made among options that include one or more of these special options, I call it a decision made in a dynamic setting.

The rule to maximize expected utility, MEU, is insufficient for decisions in dynamic settings. It does not say how to take account of information about the way in which the realization of an option would change the expected utility of the option. William Harper (1985) suggests a means of supplementing MEU to make up this deficiency. He starts with causal decision theory's version of MEU. Then he adds some of Richard Jeffrey's (1983) ideas concerning ratifiability, and some of game theory's ideas concerning mixed strategies. After presenting some decision problems that bring out the insufficiency of MEU in dynamic settings, I present Harper's proposal, point out some of its limitations, and present another proposal that avoids those limitations.

1. The Insufficiency of MEU

I agree with Harper that MEU should be given a causal interpretation. Accordingly, I assume that the formula for the expected utility for an option weights the utility for an option-state pair by the probability that the state would obtain if the option were realized.[2] I also agree with Harper that MEU should be given a ratificationistic interpretation.[3] Accordingly, I assume that MEU enjoins the realization of a ratifiable option, i.e, an option that would have maximum expected utility if it were realized.[4]

Furthermore, I make some idealizations in applications of MEU. In particular, I assume that the decision maker knows all a priori truths. This enables one to put aside models of deliberation that suppose that calculations of expected utilities supply new information.[5] Also, I

PSA 1986, Volume 1, pp. 438-449

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193144 Published online by Cambridge University Press

assume that thought is instantaneous. Under this assumption, when a decision maker reaches a decision on the basis of his evidence, he reaches it on the basis of the evidence he has at the time of the decision. Two assumptions are not among my idealizations. I do not assume that the decision maker knows what he will choose. Also, I do not assume that he knows that he will choose rationally.[6]

As a final preliminary, let us note that MEU is proposed for decision problems with a finite number of options. Consequently, I restrict our attention to such problems.[7]

The insufficiency of MEU for decisions in dynamic settings is suggested by three well known decision problems. They provide cases where MEU fails to explain our intuitions about rational choice. Admittedly, our intuitions in these decision problems are not uncontroversial. As I will see below, Harper regards some of the decision problems as pathological. Nonetheless, these intuitions are widespread so that decision rules that accommodate them are of interest.

The three decision problems have a common set-up. They involve a predictor who predicts the decision maker's choice, and prior to that choice adjusts the payoffs of the decision maker's options according to his prediction. His prediction is based on the decision maker's disposition concerning his choice. And this disposition is some psychological state of the decision maker that causes his choice.[8] The decision maker is unable to randomize his decision, say, because the world is deterministic, or because randomization is heavily penalized. Also he is aware of the payoff structure of his decision problem, and is certain of the predictor's accuracy. Hence his decision makes him certain of the payoff of his decision. This situation is essentially the same as in Newcomb's problem, a classic problem of decision in a dynamic setting.

In presenting our three decision problems, I standardize and simplify. In particular, I omit the associated stories. Each decision problem is represented abstractly by a matrix. A row represents an option, a column represents a prediction, and a matrix entry represents the payoff of an option given a prediction. The capital letter to the left of a row designates an option, and the corresponding small letter at the top of a column designates the prediction of that option. The matrix entry for the option realized, and for the prediction corresponding to it, indicates the payoff the decision maker is certain he will receive. The matrix entries for the other options given the prediction indicate the payoffs he is certain he would have received if he had realized those other options instead.

The first decision problem is the case of Death in Damascus (Gibbard and Harper 1978, Sec. 11). See Figure 1.

Death in Damascus

	a	b
A	1	10
B	9	0

Figure 1

Here A is the rational choice. It is the rational choice, although it is not ratifiable, because the option that has higher expected utility given A, namely B is also not ratifiable. This case shows that an option can be the rational choice although not recommended by MEU. Hence it shows that some principles besides MEU are necessary for a full account of rational choice.[9]

The second decision problem is the case of The Nice Demon (Skyrms 1982, p. 706). See Figure 2.

The Nice Demon

	a	b
A	1	0
B	0	2

Figure 2

Here the rational choice is B. This is so even though A is ratifiable as well as B. The case shows that not every option recommended by MEU is a rational choice. Hence it also shows that some decision principles besides MEU are needed.

The third decision problem is The Shell Game (Skyrms 1984, pp. 84-86). See Figure 3.

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193144 Published online by Cambridge University Press

The Shell Game

	a	b	c
A	0.1	0	0
B	0	10	200
C	0	100	20

Figure 3

Here the rational choice is C. It is the rational choice even though A is ratifiable and C is not. This case shows that an option can be a rational choice although contrary to the recommendation of MEU. It suggests that the decision principles supplementing MEU may occasionally override MEU.[10]

The foregoing cases of decisions in dynamic settings create the problem addressed here. How should MEU be supplemented or revised to handle such decisions?

2. Harper's Proposal

William Harper (1985, Sec. 1.2) proposes a rule for decisions in dynamic settings. As I interpret it, the rule says to realize a ratifiable option that has maximum expected utility at the last moment for deliberation.[11] In other words, it enjoins one to discard options that are not ratifiable, and from those remaining, realize one that has maximum expected utility just before its realization. In recommending a ratifiable option, the rule agrees with MEU. But in using expected utility at the last moment for deliberation to choose among ratifiable options, the rule goes beyond MEU. Let us call the rule MR since it recommends a form of maximization among the ratifiable options.[12]

It is clear that MR accommodates cases with more than one ratifiable option, such as the case of The Nice Demon. But does it accommodate cases with no ratifiable option, such as the case of Death in Damascus? The answer is, no, not strictly speaking. But as Harper shows, MR accommodates the cases obtained from such cases by introducing randomized choices, or mixed strategies. In the mixed extension of the case of Death in Damascus, there is a mixed strategy, (1/2 A, 1/2 B), that is ratifiable.[13] And since it is the only ratifiable option, MR recommends it. In the mixed extension of The Shell Game, although A is the only ratifiable pure strategy, there is a mixed strategy, (1/2 B, 1/2 C), that is also ratifiable.[14] If the probabilities of b and of c are not negligible at the last moment for deliberation, this mixed strategy has higher expected utility than A at that time, and MR recommends the mixed strategy over A.

Harper's proposal is plausible for a wide range of cases, but it has two drawbacks in the cases that interest us. First using expected utilities at the last moment for deliberation to decide among the ratifiable options does not always yield an intuitively rational choice. Suppose that in the case of The Nice Demon (Figure 2), the probability of a is high at the last moment for deliberation. Then at that time, the expected utility of A is greater than the expected utility of B. So MR recommends A.[15] This recommendation may be plausible in cases where thought takes time. But given our idealization that thought is instantaneous, A is the wrong choice, as I saw above.

The general problem with expected utilities at the last moment for deliberation is that they are out of date when a decision is reached. They do not take account of the information provided by the decision itself. A method of deciding among ratifiable options should be based upon the information one anticipates having if those options are realized.

The second drawback with MR is that it does not yield satisfactory results in cases where mixed strategies are not available. For instance, it makes no recommendation in the case of Death in Damascus (Figure 1), and the wrong recomendation in The Shell Game (Figure 3). Harper (1985, Sec. 1.2) says that such cases are pathological, and that a rule for decisions in dynamic settings need not apply to them.[16] But I hold that there are choices that are intuitively rational in these cases and that a comprehensive rule for decisions in dynamic settings ought to yield these choices.

3. Optimal Choice

The primary goal of a decision is to make an optimal choice. MEU and MR prescribe secondary goals of decision, specifically, goals for pursuit of the primary goal when one is uncertain of the choices that are optimal. As I have seen, MEU and MR are insufficient for decisions in dynamic settings. In order to obtain a comprehensive rule for decision given uncertainty, I will rebuild from the ground up. In this section I will consider the nature of an optimal choice. Then, in the final section, I will prescribe a way of attempting an optimal choice given uncertainty, even when the setting is dynamic.

Let us start by reviewing two familiar accounts of optimal choice. Both accounts have a common orientation. Both assert that an option is an optimal choice if and only if it has an outcome at least as good as the outcome of any other option. Also, both accounts take the outcome of an option to be the possible world that would obtain if the option were realized.[17] The difference between the two accounts stems from a difference in their interpretations of the conditional used to define the outcome of an option. Although both accounts take the conditional as a Stalnaker conditional, they impose different restrictions on the process of minimal revision used to obtain the nearest antecedent-world.[18]

According to the first account of optimal choice, the process of minimal revision fixes conditions prior to the time of choice, and fixes causal laws except for suspensions needed to permit adjoining options to those pre-choice conditions. I call the resulting account of optimal choice consequentialism since according to it differences in the

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193144 Published online by Cambridge University Press

outcomes of options arise because of differences in the consequences of options. It is associated with causal decision theory.

Consequentialism faces a serious problem. There are cases where the status of options is not independent of the status of features of the pre-choice context, for example, cases where an option and a prior disposition to realize the option are each optimal only if the other is. In such cases it is possible that some option is optimal although it is not optimal _given_ the pre-choice context. Consequentialism, however, ranks options strictly according to their standings _given_ the pre-choice context. Hence it errs when there is a divergence between nonconditional and conditional optimality.

To achieve independence, the second account of optimal choice imposes new restrictions on the process of minimal revision used to obtain the nearest antecedent-world. According to these restrictions, when an option is entertained, one first removes every feature of the pre-choice context such that the status of options is not independent of the status of that feature. Typically this requires removing the decision maker's disposition concerning his choice.[19] I call the time reached after backtracking in this way _the independence point_. Then one reconstructs events after the independence point so that causal laws are preserved, except for violations needed to accommodate the option entertained.

Since adjoining an option to the independence point generates a pre-choice context that goes up for evaluation together with the consequences of the option in that pre-choice context, I call the second account of optimal choice _holism_. The backtracking involved is associated with evidential decision theory. But here it is motivated by an independence condition rather than an evidential interpretation of expected utilities.

Although holism does avoid the problems concerning independence, it faces other problems. As shown by familiar arguments against evidential decision theory, if outcomes are defined in a way that involves backtracking, ranking options according to outcomes is insufficiently sensitive to the causal efficacy of options.

Since consequentialism and holism are both unsatisfactory, I need a new compromise account of optimal choice. In constructing one below, I assume that rationality provides a general goal of decision and that the nature of an optimal choice depends upon this goal of decision. Specifically, I assume that an optimal choice is what attains the goal of decision. Given these assumptions, my task is to formulate the goals of decision presumed by consequentialism and holism, obtain an appropriate compromise goal, and from it derive an account of optimal choice.

Consequentialism fixes the pre-choice context so that options are evaluated on the basis of their consequences. It presumes that the optimization of the consequences of one's choice is the goal of decision. Let us call this goal _opt-c_. Holism, on the other hand, backtracks from the pre-choice context so that options are evaluated in conjunction with earlier related parts of a possible life story, as well as with subsequent related parts of the possible life story. It presumes that the goal of decision is the optimization of one's entire life. Let us call this goal opt-l. Consequentialism and holism err because each claims that the goal it expresses is the exclusive goal of decision. As a result, they mishandle decision problems where opt-c and opt-l conflict. Holism overlooks opt-c and so overlooks the importance of causal efficacy. And consequentialism overlooks opt-l, and so overlooks the importance of independence.

In order to find an appropriate compromise between opt-c and opt-l, let us begin by considering cases that are ideal (except for the possibility of conflict between opt-c and opt-l). First, as the points about independence show, one should evaluate an option with respect to the pre-choice context obtained for it by reconstruction from the independence point. Adopting this perspective, opt-c yields for each option the goal to optimize consequences with respect to the pre-choice context constructed for it. In decision problems where the pre-choice context constructed for each option is the actual pre-choice context, the goals for all options are equivalent, and opt-c yields a nonrelative recommendation. But in other cases, these goals are not equivalent, and opt-c yields only recommendations relative to the pre-choice context for an option. Second, as the points about causal efficacy show, an option is optimal only if it meets opt-c relative to the pre-choice context for it. Thus, although opt-c is relative, it has priority over opt-l. Given that opt-c is relative and has priority, the compromise goal I seek is clear. It is to (1) restrict attention to options that meet opt-c relative to the pre-choice contexts for them, and (2) among those options, meet opt-l. Accordingly, in ideal cases an optimal choice is one that meets this hierarchical goal.

Next consider cases that are nonideal in one respect, namely for some option *o* there is an objection to pursuing opt-c with respect to *o*. Since opt-c has priority, this objection has to stem from opt-c with respect to other options. Specifically, the objection has to be that pursuit of opt-c with respect to *o* leads to other options such that pursuit of opt-c with respect to them eventually leads back to o. That is, the objection has to be that pursuit of opt-c with respect to *o* is futile from the point of view of opt-c. In cases that are nonideal in this way, I say that opt-c can be *satisfied* without being attained. I say that an option *o* *satisfies* opt-c relative to the pre-choice context for *o*, even if *o* does not attain opt-c relative to the pre-choice context for *o*, if pursuing opt-c with respect to *o* is futile from the point of view of opt-c. Taking these nonideal cases into account, the appropriate compromise goal is the hierarchical goal of meeting opt-l among the options that *satisfy* opt-c. Accordingly, I say in general that an optimal choice is one that meets this broader compromise goal.

The preceding paragraph presents the essentials of my account of optimal choice in nonideal cases. But some points need further clarification. First, let us define more precisely the *satisfaction* of opt-c by an option relative to the pre-choice context for the option. Suppose that *o* and *o*′ are options. Let us say that there is a *path* (of weak improvement) from *o*′ to *o* if and only if there is a series of options starting with *o*′ and ending with *o* such that, except for the first element, each element has consequences at least as good as its predecessor relative to the pre-choice context for its predecessor. Let us also say that an option *o* is *opposed* by an option *o*′ just in case *o*′ has better consequences than *o* relative to the pre-choice context for

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193144 Published online by Cambridge University Press

o', and there is no path from *o'* back to *o*. Then I say an option *o* *satisfies* opt-c if and only if *o* is unopposed, i.e., there is no option that opposes *o*.

Also, let me show that according to my account of optimal choice, there is in fact at least one optimal choice in every finite decision problem. To do this, it suffices to show that every finite, nonempty set of options contains at least one unopposed option. So consider a set of *n* options. Take any option of the set. Call it *o1*. *o1* cannot be opposed by itself. Suppose that *o1* is opposed by some other option. Call it *o2*. Then there is no path from *o2* to *o1*. Hence *o2* is not opposed by *o1* or itself. Suppose that *o2* is opposed by some third option. Call it *o3*. Then there is no path from *o3* to *o2* Furthermore, there is no path from *o3* to *o1*. For if there is a path from *o3* to *o1*, then there is a path from *o3* to *o2*. Hence *o3* is not opposed by *o1*, *o2*, or itself. Suppose that *o3* is opposed by some fourth option. Call it *o4*. Then by similar reasoning *o4* is not opposed by *o1*, *o2*, *o3*, or itself. Continuing in this way, I must eventually reach an unopposed option. For if *on* is unopposed by *o1*, ..., *on*, then it is unopposed. Q.E.D.

4. A New Rule for Decisions in Dynamic Settings

My objective is a rule for decisions given uncertainty, in particular, a rule that handles decisions in dynamic settings. Now that I have an account of optimal choice I am in a position to formulate such a rule. I can do it by prescribing a means of attempting an optimal choice given uncertainty. Since I accept the general idea of maximizing expected utility in the face of uncertainty, my task is merely to formulate an expected utility principle that fits my account of optimal choice.

First, I formulate expected utility principles for the accounts of optimal choice associated with opt-c and opt-1. From opt-c, I obtain, for each option *o*, the principle to realize an option whose consequences with respect to the pre-choice context for *o* have maximum expected utility. To be more precise, let *EUo*(*o'*) = ∑ *Po*(*o'*>*si*)*Uo*(*o'*,*si*), where the conditional that *o'*>*si* is interpreted as in consequentialism and the supposition that *o* is interpreted as in holism. Then, for each *o*, I obtain the principle to maximize *EUo*(*o'*). Let us call the principle schema MC. From opt-1, I obtain the principle to realize an option *o* whose outcome with respect to the independence point has maximum expected utility. To be more precise, let *EUo*(*o*) = ∑ *Po*(*o*>*si*)*U*(*o*,*si*), where the conditional that *o*>*si* is interpreted as in holism. Then I obtain the principle to maximize *EU*(*o*). Let us call this principle ML.

Next, I combine MC and ML to obtain the appropriate form of the expected utility principle for the account of optimal choice associated with our compromise goal. Since the compromise goal enjoins one to meet opt-1 among the options that satisfy opt-c, I obtain the principle to meet ML among the options that satisfy MC.[20] I call this principle HM because it recommends the *hierarchical maximization* of two kinds of expected utility.

As with opt-c, I say that an option *o* can *satisfy* MC without meeting MC relative to *o*. Specifically, I define satisfaction of MC the same

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193144 Published online by Cambridge University Press

way as satisfaction of opt-c, except substituting expected utilities for utilities. Now I say more generally that there is a path (of weak improvement) from o' to o if and only if there is a series of options starting with o' and ending with o such that, except for the first element, each element, conditional on its predecessor, has an expected utility as least as great as its predecessor's. I also say that an option o is opposed by an option o' just in case, conditional on o, o' has higher expected utility than o, i.e., EUo(o') > EUo(o), and there is no path from o' back to o. And finally I say that an option satisfies MC just in case it is not opposed by any option.

HM handles the three dynamic cases discussed in Section 1. In each of the cases, since the predictor is known to be accurate, the expected utility with respect to MC of an option o' given an option o equals the payoff of o' given that o is predicted. Also, since the predictor is known to be accurate, the expected utility of an option with respect to ML equals the payoff of the option given that it is predicted. In the case of Death in Damascus (Figure 1), each option is unopposed. So each option satisfies MC. Moreover, A has the highest expected utility with respect to ML. Therefore HM recommends A. In the case of The Nice Demon (Figure 2), each option is unopposed. So each option satisfies MC. And since B has the highest expected utility with respect to ML, HM recommends B. Finally, in The Shell Game (Figure 3), each option is unopposed and so satisfies MC. And HM recommends C since C has the highest expected utility with respect to ML.

I conclude that HM is the proper extension of MEU for decisions in dynamic settings.

Notes

[1]I am indebted to Reed Richter for stimulating correspondence on the topic of this paper.

[2]For arguments in favor of causal decision theory, and the particular version I adopt, see Allan Gibbard and William Harper (1978).

[3]For arguments in favor of ratificationism, see Richard Jeffrey (1983, Sec. 1.7) and William Harper (1985, Sec. 1.1). Jeffrey proposes ratificationism for evidential decision theory, and Harper modifies Jeffrey's proposal for causal decision theory. See also Weirich (1985) for a version of ratificationism suitable for causal decision theory.

[4]Harper's definition of a ratifiable option is slightly different. He says a ratifiable option is one that at the last moment for deliberation maximizes expected utility with respect to the condition that it is realized. However my definition is equivalent to his definition in the decision problems discussed below. There I assume that at the last moment for deliberation, for all options o and o', the decision maker knows the payoff of o' if o were realized. It follows that EUo(o') at the last moment for deliberation equals the value EU(o') would have if o were realized. (See note 11 and Section 4.)

[5]See Brian Skyrms (1982) and Ellery Eells (1984).

[6]A decision rule's job is to say what choice is rational, even in cases where the decision maker is certain that he will choose irrationally.

[7]These decision problems do not arise only from cases where the number of options is finite. They also arise from cases where the number of options is infinite but the number of *interesting* options is finite.

[8]The decision principles I discuss do not assume that free choice and causal determinism are compatible. However the recommendations of the decision principles depend on whether compatibilism is assumed in the cases to which the principles are applied. And compatibilism is assumed in all the cases I consider.

[9]See also Reed Richter (1984) for a discussion of this case.

[10]In presenting The Shell Game, Skyrms supposes initial probabilities of 1/3 for each of *A*, *B*, and *C*. In Skyrms's model for deliberation, these probabilities may indicate rational propensities to choose *A* and *B*. However, given our idealization that the decision maker knows all *a priori* truths, these probabilities indicate *irrational* propensities to choose *A* and *B*.

[11]Strictly speaking there is no last moment for deliberation. I assume, however, that the expected utility of an option converges to some value as time approaches the moment of decision, and I use that value as the value of the expected utility of the option at the last moment for deliberation.

[12]A predecessor of MR was advanced by Howard Sobel (1983, p. 166). See Harper (1985, note 3) for some reasons to modify Sobel's proposal.

[13]Assuming that the predictor cares only about making a successful prediction, (1/2 *A*, 1/2 *B*) and (5/9 *a*, 4/9 *b*) form a Nash equilibrium.

[14]Assuming that the predictor cares only about making a successful prediction, (1/2 *B*, 1/2 *C*) and (2/3 *b*, 1/3 *c*) form a Nash equilibrium.

[15]If thought takes time, then because of the time required for decision, it may be appropriate that a decision fit evidence and desires at some time earlier than the time of decision. In particular, it may be appropriate that a decision fit the evidence and desires in the total mental state that is the immediate cause of the decision.

[16]Jeffrey (1983, Sec. 1.7) takes a similar stand on cases without a unique ratifiable option.

[17]We might take outcomes as something less complex than possible worlds. But this would make accounts of outcomes more complex. Since nothing here turns on ontological simplicity, I will not explore this way of achieving it.

[18]For simplicity, I ignore issues concerning Stalnaker's conditional and the law of conditional excluded middle. None of these issues plays a critical role in what follows.

[19]It would be useful to have a fully developed theory of independence, but providing one is beyond the scope of this paper. In the absence of such a theory, I discuss only cases where intuitions about independence are fairly clear.

[20]Ellery Eells (1985, p. 477) presents a decision problem similar to The Shell Game where his intuition about the rational choice is contrary to this principle. But intuitions like those in The Shell Game confirm the principle in that case.

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193144 Published online by Cambridge University Press

References

Eells, Ellery. (1984). "Metatickles and the Dynamics of Deliberation." Theory and Decision 17: 71-95.

-------------. (1985). "Weirich on Decision Instability." Australasian Journal of Philosophy 63: 473-478.

Gibbard, Allan and Harper, William. (1978). "Counterfactuals and Two Kinds of Expected Utility." In Foundations and Applications of Decision Theory, Volume 1. Edited by C.A. Hooker, J.J. Leach, and E.F. McClennen. Dordrecht: Reidel. Pages 125-162.

Harper, William. (1985). "Ratifiability and Causal Decision Theory." PSA 1984, Volume 2. Edited by Peter D. Asquith and Philip Kitcher. East Lansing, Michigan: Philosophy of Science Association.

Jeffrey, Richard. (1983). The Logic of Decision. 2nd ed. Chicago: University of Chicago Press.

Richter, Reed. (1984). "Rationality Revisited." Australasian Journal of Philosophy 62: 392-403.

Skyrms, Brian. (1982). "Causal Decision Theory." The Journal of Philosophy 79: 695-711.

-------------. (1984). Pragmatics and Empiricism. New Haven: Yale University Press.

Sobel, J. Howard. (1983). "Expected Utilities and Rational Actions and Choices." Theoria 49: 159-183.

Weirich, Paul. (1985). "Decision Instability." Australasian Journal of Philosophy 63: 465-472.

https://doi.org/10.1086/psaprocbienmeetp.1986.1.193144 Published online by Cambridge University Press