

1

Introduction and
sketch of the main argument

1.1 Two principles of rationality

The theory of rational choice and preference, as it has been developed in the past few decades by economists and decision theorists, rests on a pair of principles. The first, the *weak ordering principle* (WO), as it is usually formulated, takes rational choice to consist in the maximization of a weak preference ordering defined over the set of feasible alternatives. Adopting the usual terminology of speaking of x as weakly preferred to y when x is either (strictly) preferred to y or x and y are indifferent, a preference relation weakly orders a feasible set X just in case it is (1) *connected* – if x and y are any two alternatives in X , then either x is weakly preferred to y or y is weakly preferred to x (possibly both), and (2) *fully transitive* – for any three alternatives x , y , and z in X , if x is weakly preferred to y , and y is weakly preferred to z , then x is weakly preferred to z . Correspondingly, choice can be said to maximize such an ordering on X when the alternative x chosen satisfies the condition that there be no other alternative y in X such that y is (strictly) preferred to x . As it turns out, however, there is a quite distinct condition that is also presupposed – albeit usually only implicitly – in the weak ordering principle, namely, that the ordering is context free. To say that an ordering is context free is to say that the ordering on any set X can be built up piecewise, by combining information concerning how the agent preferentially orders pair sets of alternatives in X when they are presented in isolation from any of the other alternatives in X .

The second principle, the *independence principle* (IND), applies to a special set of options. It requires that rational choice among risky options – those involving chance-conditioned outcomes – satisfy an independence condition with respect to component gambles or pay-offs. On one frequently employed version of this principle, the requirement is that for any three gambles g_1 , g_2 , and g_3 , g_1 is weakly preferred to g_2 iff a (compound) gamble that yields g_1 with probability p , and g_3 with probability $1 - p$, is weakly preferred to a gamble that yields g_2 with the same probability p , and g_3 with probability $1 - p$, for any $p > 0$.

It is a familiar enough result that if an agent's preference and choice behavior satisfies these two principles, as well as certain other somewhat more technical postulates, and if various structural assumptions are satisfied, then his behavior can be represented by a utility function defined over payoffs and well-defined gambles that satisfies the expected utility hypothesis.¹ That is, there exist numbers $u(g_1)$, $u(g_2)$, and so on that can be assigned to component gambles, g_1 , g_2 , and so on, such that

- (1) the magnitudes of the numbers reflect the agent's preferences, that is, $u(g_1) \geq u(g_2)$ iff g_1 is weakly preferred to g_2 ,

and

- (2) the utility to the agent of any gamble is equivalent to its *expected* utility: that is, for any gamble of the form $g = [g_1, p_1; g_2, p_2; \dots; g_n, p_n]$, such that $p_1 + p_2 + \dots + p_n = 1$, $u(g) = p_1u(g_1) + p_2u(g_2) + \dots + p_nu(g_n)$.

It is also well known that if the agent's preference and choice behavior satisfies these two principles, together with a somewhat more extended set of additional postulates and structural assumptions, then his beliefs about the likelihood of various events can be represented by a well-defined subjective probability measure over sets of mutually exclusive and jointly exhaustive events. That is, they imply that his subjective "degrees of belief" about the occurrence of such events will be representable as *point* probabilities that satisfy the usual axioms of a probability measure.² Finally, these two constructive results can be logically integrated to yield an even more powerful theorem to the effect that the agent's preferences over consequences and actions, and degrees-of belief with regard to conditioning events, can be represented respectively by a utility measure and a subjective probability measure, such that the utility of any action whose possible outcomes are conditioned by various mutually exclusive and exhaustive events is equivalent to the subjective expected utility of that action.³ In very general terms, then, the weak ordering principle and the strong independence axiom are cornerstones for the modern theory of rational choice and subjective probability.

1.2 The focus of the book

The two principles in question, together with their implications, are subject to both a descriptive and a normative interpretation. That is, they can be taken as descriptive of the preference and choice behavior

of actual agents, or as prescriptive – as specifying conditions that should be satisfied by any agent. On the usual account, the sense of “should” to which appeal is made in such a case is *nonmoral*: The notion is not that agents have a moral obligation to satisfy such conditions, but simply that these conditions can be understood as constitutive of ideally rational behavior and, hence, as setting normative standards by appeal to which the agent’s choice behavior can be criticized.

It is customary, in the study of such principles, to distinguish between logical and justificatory issues. On the one hand, there is the logical question of what theorems can be proved, given that one assumes this or that set of principles of rationality. For example, many have been preoccupied with the question of what set of principles are necessary, or sufficient, for the existence of utility indicators that satisfy the expected utility property or the existence of well-defined subjective probabilities. On the other hand, there is the quite distinct question of how such principles themselves can be defended or established when they are taken either as descriptive or as normative for preference and choice.

The focus of this book is on normative and justificatory issues. I shall be concerned with the weak ordering and independence principles as normative for, rather than descriptive of, preference and choice behavior; and I shall be preoccupied with the question of how one might justify these principles when they are so interpreted, rather than with the question of their logical implications. That is, my concern will be with the question of how one might justify the claim that a thoughtful, rational decision maker ought (in some nonmoral sense of that term) to avoid violating these two principles.

1.3 The problem of justification

When the principles in question are taken as descriptive of preference and choice behavior, whether they do in fact describe the behavior of actual agents need not raise a special methodological issue. That is, they are then ordinary empirical hypotheses subject to the usual canons of empirical verification. If, however, they are taken as normative for preference and choice, this does seem to pose a methodological problem.

Recent discussions of justification suggest that one might adopt either a foundationalist approach, and seek to ground these principles in other, even more fundamental principles (which are themselves unquestionable), or a coherentist approach, and argue that the principles are normatively valid because they explicate or codify the way in which we are in fact disposed to think about evaluating and

choosing among alternatives.⁴ Some foundationalists have even suggested that the principles themselves are intuitively secure – and thus that there is no need to appeal to even more fundamental principles. It is true, of course, that any foundationalist account will have to appeal eventually to something that is itself taken to be intuitively secure. What one hopes for, however, are starting points that command nearly unanimous acceptance, at least among thoughtful and knowledgeable researchers. Unfortunately, the above-mentioned principles do not appear to meet this test. Both have been the subject of sustained, spirited, and thoughtful questioning by a number of decision theorists.⁵ Thus, they appear to be unsuitable starting points. But this last consideration would also seem to work against any “coherentist” argument as well. The principles in question do not codify the choice behavior of competent or even expert decision makers.

1.4 An alternative approach to justification

There is, however, a quite distinct and more promising approach to the justification of these principles. It can be argued that if the agent’s preference and choice behavior fails to satisfy one or the other of these principles, it will be possible to place him in a situation in which he will choose in a pragmatically indefensible manner. More specifically, the argument is that the agent will fail to achieve his intended objective or will fail to maximize with regard to his own preferences with respect to outcomes.

The appeal to such a pragmatic test for determining what is and is not rational finds an early and unusually clear expression in Hume’s *Treatise* – particularly the passage I have taken as the epigraph for this work, in which Hume argues that rationality requires choice of means that are appropriate to the realization of one’s ends.⁶ If this is adapted to the present context, a principle of choice is valid if failure to adhere to it would result in choice of means insufficient to desired ends – in the agent pursuing his objectives less effectively than he could have under the circumstances in question.

Such an appeal to pragmatic considerations is also implicit in the way in which the independence axiom is defended in an important early article by Milton Friedman and Leonard J. Savage.⁷ They argue that the independence assumption can be secured by appeal to what they term the “sure-thing” (or dominance) principle. This principle requires choice of action x over y if, for each of a set of mutually exclusive and exhaustive conditioning events, the outcome of having chosen x is always at least as good as the outcome of having chosen y and in at least one case better. So formulated, the principle captures neatly the spirit of Hume’s requirement that choice of

Cambridge University Press

978-0-521-06391-3 - Rationality and Dynamic Choice: Foundational Explorations

Edward F. McClennen

Excerpt

[More information](#)

1.5 HAMMOND'S CONSEQUENTIALIST ARGUMENT 5

means be sufficient to desired ends: The agent who violates the sure-thing principle ends up having to settle, regardless of the turn of events, for something no more, and under certain circumstances even less, desirable than what he could have secured had he chosen differently.

The pragmatic approach also finds expression in an argument to the effect that by failure to abide by this or that principle of choice the agent can be turned into a “money pump.” The suggestion is that an entrepreneur can offer such an agent a sequence of trades, each of which, in turn, he will be disposed to accept even though, as a result of accepting all of the trades, he will end up facing the same prospect he did at the outset – except that he will be poorer, and the entrepreneur richer, by some amount of money. And since it seems possible for the entrepreneur to exploit the agent repeatedly in this fashion, the implication is that he can be “pumped” for all the money he has.⁸ The money-pump argument is related, in turn, to the “dutch-book” argument appealed to by Ramsey, de Finetti, and others in defense of the claim that rational agents should have well-defined subjective probabilities satisfying the usual constraints on probability measures.⁹ Briefly put, this argument purports to show that an agent who fails to have such well-defined subjective probabilities can be manipulated into accepting a combination of money bets such that, regardless of how events turn out, he must end up losing money – that is to say, he will end up violating the dominance principle.

1.5 Hammond's consequentialist argument

More recently, the economist Hammond has offered a somewhat related, although much more formal and complex, argument in defense of both the weak ordering and the strong independence principles.¹⁰ Hammond's argument shares with the ones mentioned in the preceding section an orientation to pragmatic considerations about outcomes or consequences, and it also shares with the money-pump argument the appeal to a dynamic, or sequential, choice framework. Indeed, in Hammond's case, this dynamic framework is explicitly developed in a very systematic manner, the thrust of the argument being that acceptance of certain principles of dynamic choice requires a commitment to the two principles of expected utility theory.

In the original version of his argument, the agent who violates one or the other of these two principles can be placed in a dynamic choice situation in which, roughly speaking, what he would now prefer that he choose at some later point in the decision tree is not what he would then prefer to choose (when he arrives at that later point). That is, violations of the ordering or independence principles with regard to preferences

Cambridge University Press

978-0-521-06391-3 - Rationality and Dynamic Choice: Foundational Explorations

Edward F. McClennen

Excerpt

[More information](#)

for outcomes and gambles over outcomes can induce a “dynamically inconsistent” shift from one temporal point to another in the agent’s preferences. The suggestion, in turn, is that such shifts cannot be squared with the assumption that the agent always chooses to maximize with respect to his preferences for consequences. More recently, Hammond has reformulated this point and argued that on a suitable axiomatic formalization of such a principle of consequences within a dynamic choice framework, the weak ordering and independence principles can be shown to follow logically, that is, to be recoverable as theorems.

I find myself unpersuaded by Hammond’s argument. I believe that if his argument is construed axiomatically (as he himself intends), it is subject to a number of objections. One can discern in his argument implicit appeal to at least three distinct rationality conditions the conjunction of which, though sufficient to yield the intended theorems, can be questioned. Among other things, it can be argued that the theorem goes through only on condition that one make a rather strong presupposition about what plans are feasible for the decision maker. Given an appropriate reformulation of the criterion of feasibility and an adjustment of the rationality conditions to this new criterion, the theorems do not go through. Hammond can, of course, insist that the feasibility presupposition is acceptable, but the problem is simply that his formal construction provides no leverage one way or the other. Moreover, even if one were to resolve this issue in his favor, it remains an open question whether the consequentialist principle to which he appeals really can ground all three of the factored conditions. My sense is that on a plausible construal of his consequentialist principle, one of the requisite conditions cannot be defended.

1.6 Pragmatism and dynamic choice

Hammond’s dynamic choice framework, however, makes possible a much more precise formulation of the pragmatic arguments mentioned in Section 1.4. In particular, it can be shown that agents who are prepared to relax either the weak ordering or the independence principle do face certain potential shifts over time in their preferences. The existence of such shifts makes the agent liable to end up implementing a plan that is strictly dominated, with respect to preferences for outcomes, by another plan that is also available to him. That is, it appears that the agent will end up failing to do as well as he could, given his own preferences for outcomes.

This point can be illustrated by a very simple example. Consider the following prospects:

$$\begin{aligned}
 g_1 &= [\$2400, 1], \\
 g_2 &= [\$2500, \frac{33}{34}; \$0, \frac{1}{34}], \\
 g_3 &= [\$2400, \frac{34}{100}; \$0, \frac{66}{100}], \\
 g_4 &= [\$2500, \frac{33}{100}; \$0, \frac{67}{100}],
 \end{aligned}$$

where $[\$2400, 1]$ is to be read as “The agent will get \$2400 with probability 1,” that is, with certainty, and $[\$x, p; \$y, 1-p]$ is to be read as “The agent will get $\$x$ with probability p , and $\$y$ with probability $1-p$.” Suppose, now, the agent prefers the prospect g_1 to g_2 , but also prefers g_4 to g_3 . In the presence of certain other seemingly uncontroversial assumptions, such a preference pattern can be shown to violate the independence principle.¹¹ Consider also a prospect $g_{3+} = [\$2401, \frac{34}{100}; \$1, \frac{66}{100}]$, in which the very same events that condition the payoffs in g_3 condition marginally larger payoffs. It is plausible to suppose that the agent in question will prefer g_4 to g_{3+} and g_{3+} to g_3 .¹² Finally, suppose that the agent is exposed to these various prospects in virtue of being confronted by the sequential decision problem shown in Figure 1.1, where squares designate choice points and circles designate chance happenings. At the first




Figure 1.1




Figure 1.2

choice point, the agent faces the option of accepting the prospect $g_3 +$ outright or continuing on and (possibly) exercising a second choice. If he takes the second option, planning, in the event that circumstances permit it, to choose g_2 rather than g_1 , he exposes himself to a prospect that is equivalent (in terms of monetary payoffs and probabilities) to g_4 . To see this equivalence, note that one can mentally erase part of the problem as given in Figure 1.1 and think of the agent as facing the problem given in Figure 1.2. That is, since the agent plans to choose g_2 , if and when he reaches the second choice point, the dashed portion of Figure 1.2, designating what would have been a possible subsequent choice of g_1 , is not relevant.¹³ But in this case, the agent really just faces the problem given in Figure 1.3.

Moreover, by appeal to the usual rules for combining independent probabilities, the upper branch of this problem is essentially equivalent to the branch given in Figure 1.4, that is g_4 , which, it will be recalled, the agent prefers to $g_3 +$. Correspondingly, for an agent who conditionally plans (were he to reject the prospect of $g_3 +$ at the first choice point) to reject the prospect of g_1 (if and when he arrives at the

1.6 PRAGMATISM AND DYNAMIC CHOICE




Figure 1.3




Figure 1.4

second choice point), the problem presented in Figure 1.1 is essentially equivalent to the one given in Figure 1.5.

Moreover, since it is plausible to suppose that he will choose g_4 over $g_3 +$ in the problem given in Figure 1.5, we may infer that for the problem given in Figure 1.1, he will choose to reject $g_3 +$ and head toward the second choice point.

Cambridge University Press

978-0-521-06391-3 - Rationality and Dynamic Choice: Foundational Explorations

Edward F. McClennen

Excerpt


[More information](#)

Figure 1.5

Now suppose the agent does just this. If chance operates so that he does not arrive at the second choice point (event $-E$ takes place), he gets \$0 – which is less than the \$1 he would have received, *under the very same chance condition*, had he accepted g_3+ at the outset. Suppose, however, that chance favors him (event E takes place), and he arrives at the second choice point. Despite the plan he formulated to himself at the first choice point, what he now faces is a choice between prospect g_1 and prospect g_2 . That is, in terms of the “dashed-line” notation employed earlier, what he now faces is simply the “undotted” portion of the problem given in Figure 1.6. But *by hypothesis*, he prefers g_1 to g_2 , and thus it is plausible to suppose that he will, in fact, end up choosing g_1 . Notice, however, that what he then gets is \$2400, which is less than the \$2401 he would have received, *under the very same chance condition* (event E takes place), had he accepted g_3+ at the outset. Thus, he loses either way. That is, no matter how chance works (E or $-E$), he does worse than he would have done had he initially chosen g_3+ .

Here, then, is an example of a carefully formulated pragmatic argument. Moreover, given the continuing controversy over the status of the weak ordering and the independence principles as normative for choice, it would seem to mark out a very promising justificatory approach. If anything is likely to command universal assent among thoughtful people, it is just the notion that effective choice of means to ends is (at the very least) *a* criterion of rationality. Any argument that succeeded in grounding the weak ordering and independence