Introduction

A typical exercise in a textbook on rational decisions imagines a diner choosing a meal from the menu at a restaurant. For simplicity, it supposes that the diner ranks items on the menu using only her tastes at the moment. However, a reasonable diner considers her decision problem's spatiotemporal context. She may consider the provenance of the restaurant's food and so look toward the past. She may also consider the consequences of her choice for causes she supports, such as local farming, and so look toward the future.

To make sure that a decision's evaluation considers all relevant factors, not just events here and now, some principles of evaluation target the *possible world* that results given the decision's realization.[1] The possible world covers the decision's past, present, and future. It has spatiotemporal structure and includes all the events that occur given the decision's realization. Evaluating it includes the consideration of every relevant pro or con no matter how distant in space or time.

Comprehensively evaluating a decision by evaluating the decision's world is safe but has an obvious drawback – complexity. It makes appraisal of ordering a dish from a menu an enormously complex appraisal of the possible world issuing from the order. The possible world has so much detail that no description, no matter how elaborate, comes close to being thorough. Practicality demands simplification.

This book shows how to narrow the scope of a decision's evaluation. An evaluation of a choice from a menu need not survey all events in the decision's world, including events in the remote past, such as the burning of the library of Alexandria in 48 B.C. Although the diner cares about that calamity, evaluation of ordering a ragout tonight may ignore it.

[1] I italicize the first occurrence of a technical term in a passage that introduces it. Usage of technical terms in decision theory is not uniform, and I introduce technical terms to suit the principles I formulate with them. Each principle's formulation reviews the points motivating its terminology.

An efficient evaluation trims three types of irrelevant consideration. It puts aside, first, matters of indifference; second, nonbasic sources of utility; and, third, unavoidable events, that is, events that occur given every option available. The first and second steps narrow considerations enough to make them describable. The third trims the past, for example. The evaluation's elaboration occupies most of the book because decision theory needs new resources to justify principles that ignore sources of utility.

People manage to focus on information relevant to tasks at hand, although their methods are sometimes opaque. For example, people recognize faces more easily than do computers because people pick out relevant facial features in ways not yet formalized. Although people generally focus on relevant considerations when evaluating options, their methods lack precise description and justification. The following chapters explicate and justify some common methods of identifying relevant features of options in a decision problem; in particular, methods that focus on an option's future, its spatiotemporal region of influence, or its consequences. Rather than offer new practical advice, they make precise familiar methods of simplifying choices and give them firm philosophical foundations.

Principles

Normative decision theory formulates principles for evaluating the options in a decision problem and for identifying rational options. Taking the theory's perspective, a feature of an option is relevant to its evaluation if the feature has a role in identifying rational options. Rationality as it applies to options is a source of standards for a decision among options. Some texts adopt technical definitions of rationality according to which, for example, an option is rational if and only if it is at the top of a preference ranking of options. Such technical definitions suit certain projects, but this book treats rationality in its ordinary evaluative sense and so grounds principles of rationality in normative arguments rather than technical definitions. Its arguments draw on judgments about rationality in particular cases and judgments about rationality's general features.

An adequate representation of an agent's decision problem specifies an exhaustive set of options available at the time the decision problem arises. Comparisons of options in a decision problem establish which options are better than others, and which are best, according to the agent's goals. These comparisons ground rational decisions. For simplicity, assume that an assignment of utilities to options settles the preference ranking of the options in a

decision problem. Then, the key to a rational decision is the assignment of utilities to options.[2]

The principle of utility maximization states that a rational decision maximizes utility. This means that the decision's utility is at least as great as any alternative decision's utility. The principle's restriction to decision problems in which some decision maximizes utility ensures the possibility of compliance. Its restriction to ideal agents in ideal circumstances makes compliance feasible. The principle is too demanding for real agents with cognitive limits and time pressure.

The principle of utility maximization compares options by assigning a *utility* to each option. I take an option's utility to indicate the agent's strength of desire to realize the option. Strength of desire depends on the agent's information. If she believes that an option will have only good consequences, she desires to realize it to some extent; but if she learns that it will have only bad consequences, she then becomes averse to realizing it. Because the principle that a rational choice maximizes utility takes rationality in its ordinary evaluative sense and defines utility in terms of desires rather than choices, it is a normative principle rather than a definitional truth.

By taking utility as strength of desire, so that an option's utility is not relative to a set of preferences, I adopt a *realist* view of utility that contrasts with a *constructivist* or *interpretativist* view that extracts utility from a set of preferences. For example, the theory of *revealed preference*, common in economics, takes a constructivist view of utility. It uses utility just to represent preferences. The book's methods of simplifying choices work given a constructivist conception of utility; however, adopting a realist conception of utility strengthens normative principles of utility. Chapter 2 reviews the reasons for the realist view.

The goal of blending probability and utility in principles of choice suggests taking both probability and utility, and not just probability, to attach to a

[2] The assumption of quantitative evaluations of options as a basis for a decision is not significantly stronger than the assumption of comparisons of options. Although identifying a top-ranked option relies directly on a ranking of options and only indirectly on quantitative evaluations of options, ranking probability mixtures of options requires quantitative evaluations of the options mixed. Granting that probability mixtures of options are options, ranking options requires quantitative evaluations of options. In any case, assuming quantitative evaluations of options does not significantly affect conclusions about efficiency in deliberations. Because deliberations use options' evaluations to compare options, reasons to ignore considerations when evaluating an option (such as considerations that apply to all options) resemble reasons to ignore considerations when comparing two options (such as considerations that apply to both options). Identification of relevant considerations is similar for evaluation of options and for comparison of options.

proposition, the content of a declarative sentence. Desire may be taken as an attitude toward a proposition. A sentence stating a proposition that an agent wants to be true expresses a target of the agent's desires. For example, an agent may desire that peace prevail. Then, an object of the agent's desires is the proposition that peace prevail. Desiring that an event occur is equivalent to wanting that the event occur. This sense of desire, which Schroeder (2004) expounds, is common in philosophy, and Chapter 2 introduces it more thoroughly. Taking desire as a propositional attitude, an option's utility attaches to a propositional representation of the option.

Besides depending on information, an option's utility depends on the scope of the option's evaluation. It may consider, for example, just the logical consequences of the option's realization, the total consequences of its realization, or the possible world the option realizes. To be comprehensive, ordinary utility evaluates the world that would be realized if the option were realized. When that world is uncertain, the option's utility equals a probability-weighted average of the worlds that might be realized if the option were realized.

This book adopts *causal decision theory's* analysis of an option's utility without reviewing the arguments for it.[3] According to one version of this analysis, namely, *world-Bayesianism* as Sobel (1994: chap. 1) presents it, a decision's utility is its world's utility, or if that is unknown, a probability-weighted average of the worlds that, epistemically speaking, might be its world. A decision's world is the possible world that *would* be realized if the decision *were* made. According to a standard account of conditionals, this possible world, among all worlds containing the decision, is nearest to the decision problem's world, assuming, for simplicity, that a unique world is nearest. Introducing the decision's world using a conditional in the subjunctive mood signals that causal relations settle the nearest world with the decision's realization, as Chapter 2 explains.

An initial simplification of a decision's world puts aside matters of indifference. It takes a decision's world as a proposition specifying for each desire and aversion whether it is realized. A further simplification excises other irrelevant parts of the decision's world, for example, the decision's past. These simplifications reduce an evaluation's information-processing demands.

World-Bayesianism advances a standard of evaluation for options, not a procedure for deliberation. An outsider whose cognitive powers differ from the agent's cognitive powers may apply the standard of evaluation. Efficient

[3] Weirich (2012) reviews the arguments.

principles for evaluation of options, even taken as standards of evaluation rather than as procedures for deliberation, ask too much of cognitively limited agents, especially in adverse circumstances, but make appropriate demands of ideal agents in ideal circumstances. Ideally situated ideal agents may convert principles of evaluation into procedures of deliberation and therefore lack excuses for failing to satisfy the principles of evaluation. Although I advance principles of utility as standards of evaluation because I restrict the principles to ideally situated ideal agents, the principles also yield procedures of deliberation.

Utility analysis, because it investigates rationality for an ideally situated ideal agent, treats a normative *model*. The model is a possible world like the actual world except for including an ideal agent in an ideal decision problem. Utility analysis fills out this sketch of the model by specifying respects in which an agent and a decision problem are ideal. The model is not mathematical, although some principles of utility are mathematical. The principles, which evaluate options to identify rational options, are not part of the model because, being necessary truths, they govern all possible worlds without distinguishing the possible world the model presents. Consequently, the model is normative because of its function rather than because it has normative components. The model does not trivialize the arguments for the principles of utility, although it removes objections to the principles. Ideal agents in ideal decision problems can violate the principles, although the arguments maintain that these agents, if rational, observe the principles.

Although realism prompts an interest in efficiency, methods of achieving efficiency may incorporate some idealizations. A method may reduce information-processing demands while idealizations accompanying it ensure enough cognitive power to meet its remaining demands.

The book's efficient decision principles target ideal agents who calculate effortlessly and know all logical truths. After evaluating options, such agents can compare options and decide without any cognitive costs. However, they can only up to a limit costlessly process considerations to evaluate an option. Efficient evaluation reduces the number of considerations to process so that evaluating an option does not exceed their limit. Because such agents, although cognitively powerful, beyond their limit pay evaluation costs, an efficient method of evaluating options benefits them. Its availability removes an excuse for failing to maximize utility and justifies holding them accountable for utility maximization.

Selective adoption of idealizations assists progress toward a theory of rationality for real agents. Idealizations simplify principles of rational

choice, and future research may remove the idealizations to generalize the principles. Weirich (2004) shows how to remove, from a normative model, idealizations such as the assumption that agents have precise probability and utility assignments.[4] Until theorists generalize idealized principles so that they also cover nonideal cases, the principles are just helpful guides in real-life decisions. The idealized principles express a goal that rational agents in real cases aspire to attain, they direct consultants and computers offering aid with decisions, and they apply approximately to some real cases. Also, some implications of the principles hold robustly in the absence of idealizations. For example, in a typical decision problem, a person evaluating an option does well to put aside past events and examine only the option's consequences, even if the shortcut's precise justification depends on idealizations. In particular, the advice is apt for a diner ordering from a menu, even if its incorporation by the directive to maximize utility rests on idealizations.

Efficiency also has theoretical interest apart from its practical interest. Identifying the considerations relevant to an option's evaluation advances a theory of rationality by refining explanations of a decision's rationality. An explanation improves by using only considerations that bear on a decision's rationality. A completely satisfactory explanation, for economy, ignores irrelevant considerations. A good explanation of the rationality of a diner's decision from a menu skips fires in ancient Alexandria, for instance. Efficient decision principles ground good explanations of a decision's rationality.

Efficiency

Although explaining efficiency in evaluation of options does not require a thorough account of efficiency, a brief account clarifies the project. This section defines the type of efficiency the principles of utility analysis achieve. It paints the background for the claim that evaluating options according to their futures is more efficient than evaluating them according to their worlds.

Many types of efficiency are valuable. I examine *efficiency in doing a task*. Specifying this type of efficiency involves specifying the task – including the

[4] A model with idealizations may offer a partial explanation of a target phenomenon. Consider the explanation of a person's checkbook balance. Suppose that a model contains checks and deposits but omits fees the bank levies. The model offers a partial explanation of the balance by displaying the operation of some factors affecting the target phenomenon. A full explanation adds the fees that the model omits. A model of a decision's rationality may similarly offer a partial explanation of its rationality. A full explanation treats factors that idealizations put aside.

resources available, the working conditions, and the final product. It also involves specifying the measure of efficiency, which may target just a part of overall efficiency.

A method of doing a task may be efficient without being maximally efficient; an efficient process need not be the most efficient process. In a classificatory sense, an efficient process completes the task without wasting resources; it uses close to the minimum required. Comparative efficiency ranks methods; some are more efficient than are others. Comparisons are transitive and in finite cases identify the most efficient processes. A method of performing a task achieves maximum efficiency if and only if no alternative method is more efficient.

Performing a task efficiently may differ from performing it in a good way because efficiency, although good, differs from overall goodness. A method of production may be efficient but not good because its product is worthless. Also, a process may be efficient in some respect even if it is not efficient overall; it may achieve one type of efficiency at the expense of another more important type of efficiency. Overall efficiency weighs and combines types of efficiency to find a suitable balance among them. To simplify, a study of efficiency may treat just one component of overall efficiency.

Efficiency in a task obtains relative to a measure of performance. For travel from Chicago to New York, a driver may evaluate routes according to the number of miles or the number of hours behind the wheel. A factory may measure efficiency in an automobile's production using a combination of raw materials, time, and costs (perhaps including environmental impact), or it may use profit. In computing, performing the task with the fewest steps or in the shortest time is a convenient measure of efficiency. Another measure may use the length of the program for the task. In modal logical, a measure of efficiency for an axiomatization may use the number of symbols in the axiomatization. An alternative measure may use the lengths of proofs of a representative set of theorems using the axiomatization and a set of rules of proof.

Efficiency in a task depends on a precise specification of the task. The task may be general, such as building automobiles, or it may be a specific instance of a general task, such as building a particular automobile. Specifying the resources, conditions, and product is a way of specifying the task. Resources for building automobiles may include not only metal used to make a chassis but also workers and tools. Specifying the task also includes specifying the product. It may be an automobile that meets high standards of quality. Efficiency evaluates a process going from a specific starting point to a specific end point in a specific environment.

These points about efficiency in a task carry over to efficiency in evaluation of an option. The task is evaluation of the option for the purpose of choosing an option. The resources include an ideal agent's cognitive abilities; an identification of basic intrinsic desires and aversions, or *basic intrinsic attitudes*, that the option's world realizes; and their realizations' temporal, spatiotemporal, and dependency relations to the option. The conditions include the decision problem's being ideal so that probabilities and utilities of the option's possible outcomes are available. The product is an accurate evaluation of the option, namely, the option's utility. The measure of efficiency is the number of basic intrinsic attitudes reviewed.

Identifying the general method of evaluation that processes the fewest considerations is just a step toward overall efficiency. Simplifying choices requires many types of efficiency, and calculations of overall efficiency weigh processing just relevant considerations along with other types of efficiency. Additional types of efficiency may, for example, count the calculations an evaluation performs.

Because deliberation in a decision problem includes, besides evaluation of options, identification and selection of options, efficient evaluation does not entail efficient deliberation. Efficient deliberation may not use efficient evaluation if identifying relevant considerations is costly. Deliberation that sorts considerations by relevance may be inefficient. However, given the sorting as a resource, an evaluation that processes only the relevant considerations is efficient.

Given the uncertainty of an option's world, an option's evaluation uses probabilities of worlds that might be the option's world. In a broad sense, relevant considerations include features of evidence that influence these probabilities. However, given the probabilities as a resource, an efficient evaluation does not process the evidential considerations. An efficient evaluation, granting my specification of the task, processes only, and not necessarily all, considerations relevant to evaluating the worlds that might be the option's world.

An efficient principle of utility maximization minimizes a suitable utility assignment's demands on an agent's information about options; it efficiently obtains options' utilities. Its utility assignment uses only information necessary for generally applying the principle of utility maximization. Narrowing the scope of an option's evaluation achieves this efficiency. The narrower the scope, the fewer considerations the evaluation reviews and hence the lower its cognitive cost, other things being equal.

Efficient general principles for decisions use a minimum of the agent's information about options sufficient in every decision problem for resolving

the problem. Specialized principles may further minimize informational demands because the demands vary from problem to problem. Suppose that in a decision problem, an option realizes some aversions and no desires, and a second option realizes some desires and no aversions. The second option is superior to the first option no matter which aversions the first option realizes and which desires the second option realizes. In this case, comparison of the options need not review the two options' consequences. Assuming that the rational options a decision problem's resolution identifies are at the top of a preference ranking of options, an efficient principle for the decision problem may compare the two options without placing each option on a utility scale according to its consequences. Evaluation by consequences, although generally efficient, is not in every decision problem a maximally efficient way of ranking options.

For perspective, consider an efficient way to select three green marbles from a bag. A straightforward method draws marbles one by one, inspects each, and keeps it if it is green, until collecting three green marbles. Another method takes three marbles at once and, if necessary, replaces marbles not green to obtain three green marbles. The second method is more efficient if all the marbles in the bag are green. It takes fewer steps handling individual marbles in that case. Efficiency depends on circumstances.

Efficient general methods of evaluation make informational demands no greater than other general methods of evaluation. In special cases, some decision principles may be more efficient with respect to information processing than an efficient general principle. In case-by-case comparisons, one decision principle may not be uniformly more efficient than another but may as a general principle nonetheless be more efficient than the other principle because its general application requires less information than does the other principle's general application. This book refines for efficiency the principle of utility maximization. Although the refined principle is not efficient in every decision problem, it achieves efficiency with respect to other general decision principles. It uses an efficient general method of assigning utilities to the options in a decision problem to rank options and thereby resolve the decision problem.

Characterizing the relevant considerations for comparing options is a theoretical project. The characterization also has practical value because it licenses ignoring parts of an option's world to make choice efficient. However, theoretically efficient and practically efficient evaluations may differ. In some cases, theoretically efficient evaluations do not achieve practical efficiency because identifying relevant considerations requires reviewing all considerations anyway. Applying a method of evaluation

that is optimal in theory may not be optimal in practice because the method processes hard-to-identify considerations. A theoretically efficient evaluation procedure may not be practically efficient because of its application's cognitive demands.[5]

To illustrate this point about efficiency, consider an example from computer science. Sometimes efficient methods of computation are hard to implement. To sort n items, an efficient method uses $n \times \log(n)$ steps. To sort 1000 items, it takes $1000 \times \log(1000)$ steps, or 3000 steps. The efficient method is hard to implement. So, one might use an easy implementation of an inefficient process that uses n-squared steps. Although it takes 100,000 steps to sort 1000 items, it may be more practically efficient than its theoretically efficient rival. Similarly, an efficient decision method may consider only options' consequences, although an easy-to-implement, inefficient method may consider options' worlds. Theoretical efficiency yields practical efficiency only given idealizations that facilitate implementation.

In some cases, a theoretically simple evaluation of an option is impractical because the evaluation must review every consideration to ascertain whether it is relevant. In other cases, however, the evaluation may dismiss irrelevant considerations in batches so that it is practical. For example, it may dismiss the past without reviewing every past event. Although theoretically efficient evaluations are sometimes cognitively demanding, they have practical value in cases where people can perform them.

Some reasonable simplifications of evaluations risk failure to identify a utility-maximizing choice. Deliberation costs may justify evaluating options with respect to a single objective, such as money, although the simplification neglects relevant considerations. Reasonable simplification continues beyond the least taxing evaluations that identify a utility-maximizing choice. It trades accuracy in resolving decision problems for reduced cognitive costs. This book's assessment of decision methods for efficiency ignores such trade-offs and compares only accurate methods of resolving decision problems.[6]

[5] Shanahan (2009) states artificial intelligence's frame problem in its general form: "How do we account for our apparent ability to make decisions on the basis only of what is relevant to an ongoing situation without having explicitly to consider all that is not relevant?" Overall efficiency for deliberations has the goal of identifying relevant considerations without reviewing all considerations and so the goal of solving a philosophical version of the frame problem.

[6] Gigerenzer (2000: chap. 8) describes decision heuristics that promote efficiency in realistic environments, but at the cost of accuracy in general applications. For example, the heuristic to "Take the Best and Forget the Rest" makes a decision using one type of reason and ignoring other types of reason. Although it efficiently yields good decisions in a congenial environment, it fails when the reasons ignored contravene the reasons considered. Utility analysis advances precise general methods of