

Choosing Strategies

Suppose I need to decide whether to go off to fight for a cause in which I deeply believe, or stay home with a family that needs me and that I deeply love. What should I do? My friends say I should determine the possible outcomes of the two proposed courses of action, assign probabilities and numerical utilities to each possibility, multiply through, and then choose whichever alternative has the highest number.

My friends are wrong. Their proposal would be plausible in games of chance where information on probabilities and monetarily denominated utilities is readily available. In this case, however, I can only guess at the possible outcomes of either course of action. Nor do I know their probabilities. Nor do I know how to gauge their utilities. The strategy of maximizing expected utility is out of the question, for employing it requires information that I do not have.

Nevertheless, my friends have not given up trying to help, and so they point out that I could simulate the process of maximizing expected utility by assuming a set of possible outcomes, estimating their probabilities, and then making educated guesses about how much utility they would have. I could indeed do this, but I decide not to, for it occurs to me that I have no reason to trust the formula for maximizing expected utility when I have nothing but question marks to plug into it. Better strategies are available.

1. Two Kinds of Strategies

This section distinguishes between optimizing and satisficing strategies, and between moderate and immoderate preferences. The following three sections discuss, in turn, when satisficing strategies are rational, when they are not, and when cultivating moderate preferences is rational. Later sections offer a way of characterizing rational choice in situations where the agent's alternatives are incommensurable.

In the simplest context, one has a set of alternatives clearly ranked in terms of their utility as means to one's ends. If one is an *optimizer*, one chooses an alternative that ranks at least as high as any other. In contrast, if one is a *satisficer*, one settles for

any alternative one considers satisfactory. In this static context, though, it is hard to see the point of choosing a suboptimal alternative, even if it is satisfactory.

In a more dynamic and more typical context, we are not presented with a set of nicely ranked alternatives. Instead, we have to look for them, judging their utility as we go. In this context, optimizing involves terminating one's search for alternatives upon concluding that one has found the best available alternative. However, although optimizing involves selecting what one judges is best, it need not involve judging what is best "all things considered," because sophisticated optimizers recognize that considering all things is not always worth the cost. There may be constraints (temporal, financial, and so on) on how much searching one can afford to do. A person who stops the search upon concluding that prolonging the search is not worth the cost is also employing an optimizing strategy, albeit one of a more subtle variety.¹

Satisficing, in contrast, involves terminating the search for alternatives upon concluding that one has identified a satisfactory alternative. What distinguishes satisficing from optimizing in the dynamic context is that the two strategies employ different *stopping rules*.² Thus, if options emerge serially, a subtle optimizer might choose a known option in preference to the alternative, namely, searching for something better with no guarantee of ever finding it. The difference between satisficing and this more subtle kind of optimizing has to do with what the two strategies take into account in reaching a stopping point. At any point in the search, we may let the expected utility of stopping the search equal U , the utility of the best option we have turned up so far. The expected utility of continued search equals the probability of finding a better option, $P(fbo)$, multiplied by the utility of finding a better option, $U(fbo)$, minus the cost of further search, $C(fs)$. At some point, the satisficer stops because he believes U is good enough. In contrast, the subtle optimizer stops because she believes that $P(fbo)U(fbo) - C(fs)$ is less than zero. Even if the two stopping rules happen to converge on the same stopping point, they do so for different reasons and require different information.³

Unlike the optimizer, who stops searching when she either has considered all her options or has run up against things like time constraints, the satisficer stops the

1. Stocker (1990, 311–316) argues that optimizing, even in this subtle sense, is morally and rationally problematic.

2. Herbert Simon's (1955) idea is that, given our limited capacity to acquire and process information, we economize on our limited capacity by setting a concrete goal and then reasoning back to conclusions about what course of action would achieve that goal. This is what Simon means by satisficing. The notion of satisficing as a stopping rule is implied, and later becomes explicit (in Simon 1979, 3). Simon treats satisficing as a surrogate for optimizing under particular information constraints and, so far as I know, treats our limited information as an external constraint. I think such constraints are often more accurately viewed as being partly self-imposed. Within the context of constraints that are in part self-imposed, the distinction between satisficing and optimizing becomes more interesting, as explained in section 2.

3. When pressed to justify stopping a search upon finding a satisfactory alternative, people may say further search would not have been worth the cost. Does that mean they were optimizing after all? Only if the cost of further search was what in fact made them stop. If the circumstance that actually stops a search is the finding of a satisfactory alternative, then it is a case of satisficing, no matter what people say after the fact in defense of the choice. One can defend two acts in the same way without implying that the two acts are the same kind of act. For more on how the two stopping rules differ and on how they might usefully be combined, see the appendix to this chapter.

search upon identifying an alternative as good enough.⁴ For example, suppose you enter a cafeteria seeking a nutritionally balanced and reasonably tasty meal. You then proceed down the cafeteria line surveying the alternatives. If you are satisficing, you take the first meal that you deem nutritionally and aesthetically adequate. If you are optimizing, you continue down the line surveying alternatives until you reach the end of the line or run out of time. You then take the meal you consider optimal, either in comparison to the other known options or in comparison to the alternative of further search. A satisfactory meal may or may not be optimal. Likewise, as cafeteria patrons know only too well, the best available meal may or may not be satisfactory. Of course, if you switch from one stopping rule to the other, you might end up choosing the same meal, but you will be choosing it for a different reason. Therefore, neither rule is reducible to the other.

Neither can satisficing be equated with the more subtle kind of optimizing that takes the cost of searching for more-than-satisfactory alternatives into account. Satisficers select the satisfactory alternative because they find it satisfactory, not because they find that stopping the search at that point would maximize utility.

With this characterization of satisficing in mind, we can now clarify the difference between satisficing and *moderation*. Satisficing contrasts with optimizing. Being moderate, however, contrasts not with optimizing but with being immoderate. Being an optimizer does not entail being immoderate, and being a satisficer does not entail that one would be satisfied with a moderate bundle of goods. A person could be both a moderate and an optimizer, for the maximally satisfying bundle of goods for a given person may well be of moderate size. Likewise, a person could be both a satisficer and an immoderate, for a given satisficer may have wildly immoderate ideas about what counts as satisfactory. Consider a person whose goal in life is to be a millionaire—not a billionaire, mind you, just a millionaire—by the age of thirty.

2. When Satisficing Is Rational

There is an apparent incongruence between the theory and practice of rational choice. Theory models rational choice as optimizing choice, yet in practice, satisficing is ubiquitous. We could explain the incongruence away by saying that when people think, they seek something satisfactory, what they really seek is something optimal. But satisficing can be reconstructed as a subtle kind of optimizing strategy only on pain of attributing to people calculations they often do not perform (and do not have the information to perform) and intentions they often do not have. This section tries to explain satisficing in terms of thought processes we can recognize in ourselves.

4. There may be no precise way to characterize “good enough”. But what people consider good enough seems relative to expectations. As expectations rise, the standards by which an option is judged good enough also tend to rise. This fact can be tragic. It can rob people of the ability to appreciate how well their lives are going, all things considered. Of course, it is rational to set *goals* with an eye to what is attainable, raising one’s sights as higher goals become attainable. But raising the standard by which we deem our situation satisfactory is harder to fathom. Perhaps people are psychologically incapable of aiming at higher goals without simultaneously reformulating their notions of what is satisfactory. I do not know.

Satisficing will emerge as a real alternative to optimizing, and thus as a strategy that can be evaluated, criticized, and sometimes redeemed as rational.

We begin with the observation that people have a multiplicity of goals. For example, a person can desire to be healthy, to have a successful career, to be a good parent, and so forth. Some goals are broad and others narrow, relatively speaking. Further, a given goal might be encompassed by another, in the sense that the narrower goal's point—the reason for it being a goal—is that it is part of what one does in pursuit of a larger goal. For example, Kate might want to upgrade her wardrobe because she cares about her appearance because she wants a promotion because she cares about her career. Suppose she believes that achieving her various goals is instrumental to or constitutive of achieving a broader goal of making her life as a whole go well. To mark the difference in breadth between Kate's concern for her life as a whole and her concern for particular aspects of her life (such as her health or her career), let us say Kate seeks a *local* optimum when she seeks to make a certain aspect of her life go as well as possible. Kate seeks a *global* optimum when she seeks to make her life as a whole go as well as possible.⁵

Optima can be defined as such only within the context of the constraints under which goals are pursued. (Thus, when economists speak of maximizing utility, it goes without saying that they are talking about maximizing utility subject to a budget constraint.) We pursue goals subject to the limits of our knowledge, time, energy, ability, income, and so on. More intriguing, however, is that we typically operate under additional constraints that we have deliberately imposed upon ourselves, as if the constraints imposed on us by external circumstances were too loose. For example, if Tom spends an evening at a casino, he is externally constrained (by his savings and his borrowing power) to spend no more than, say, \$50,000. What actually defines his set of options over the course of the evening, though, is the \$100 budget constraint that he has *chosen* to impose on himself.⁶

To give another example, in fleshing out the task of buying a house, we need to make some prior decisions. We decide how long to look, how much money to spend, what neighborhoods to consider. We knock only on doors of houses displaying “for sale” signs rather than on every door in the neighborhood. To some extent, these constraints are imposed on us by mundane external factors, but they also have a striking normative aspect, for they are in part rules of conduct we impose on ourselves; we take it upon ourselves to make our constraints more precise and more limiting so as to make our choice set more definite. Local optimizing would often be neurotic and even stupid if local goals were not pursued within compartments partly defined by

5. I borrowed the terms “local” and “global” from Elster (1984, 9), although a rereading of his text reveals that the way he uses the terms bears little resemblance to the way they are used here. (Elster says the definitive difference between locally and globally maximizing machines is that the latter, unlike the former, are capable of *waiting* and *indirect* strategies.)

6. For someone wanting to construct a tractable mathematical model, it may be easier to ignore the felt experience of pursuing local goals under self-imposed constraints and concentrate instead on the global perspective, from which self-imposed constraints appear, more or less, as preferences about how to operate within external constraints. I want to explain satisficing in terms of thought processes we can recognize within ourselves, though. So for my purposes, the fact that we have both broader and narrower perspective cannot be ignored.

self-imposed constraints. The constraints we impose on our narrower pursuits can keep narrower pursuits from ruining the larger plans of which they are part.⁷

If we look at life as a whole, we see that life as a whole will go better if we spend most of it pursuing goals that are narrower than the goal of making life as a whole go better. That is why it is rational to formulate and pursue local goals. But it also is rational to prevent narrower pursuits from consuming more resources than is warranted by the importance (from the global perspective) of achieving those narrower goals. Accordingly, when we pursue narrower goals, we pursue them under self-imposed constraints.

Although the constraints we impose on ourselves are imposed from a more encompassing perspective, it is only within the narrower perspective that we become subject to self-imposed constraints. (Of course, we are subject to external constraints, limited incomes and such, from any perspective.) Self-imposed constraints can be applied only *to* narrower pursuits and can be applied only *from* the perspective of a more encompassing pursuit. In more familiar terms, the point is that, because we have broader objectives, there are limits to what we will do for the sake of our wardrobe, or for the sake of a promotion, or for the sake of a career.

Having distinguished between local and global optimization, we can now explain when satisficing is rational. Michael Slote believes the optimizing tendency can be self-defeating. He says, “A person bent on eking out the most good he can in any given situation will take pains and suffer anxieties that a more casual individual will avoid.” And he asks us to consider “how much more planful and self-conscious the continual optimizer must be in comparison with the satisficer who does not always aim for the best and who sometimes rejects the best or better for the good enough.”⁸ In short, that one has an opportunity to pursue the good is not by itself a compelling reason to pursue the good. Surely, Slote has an important point. Just as surely, however, his point applies to local optimizing rather than to optimizing as such. From the global perspective, seeking local optima can be a waste of time. Global optimizers seek local optima only when doing so serves their purposes. For that reason, satisficing is a big part of a global optimizer’s daily routine. A compulsive seeking of local optima may go with being immoderate, but not with being a global optimizer. Effort can have diminishing returns, so a global optimizer will be careful not to try too hard. Local optimizing often gives way to satisficing for the sake of global optimality.

From the global optimizer’s point of view, the process of buying a house provides a good example of how satisficing can be rational. When we choose a house, we might proceed by seeking the best available house within constraints—within a one-month time limit, for example. We impose such a limit because we have goals other than living in a nice house, and finding a house competes with them for our time and energy. Or we might look for a satisfactory house and cease looking when

7. As Jules Coleman has pointed out to me, what Gauthier (1986, 170) calls constrained maximization is a particularly interesting kind of local optimizing under self-imposed constraints. Constrained maximizers seek maximum payoffs in prisoner’s dilemmas, subject to this constraint: they will cooperate (and thus pass up the opportunity to unilaterally defect) if the expected payoff of cooperating is higher than the known payoff of *mutual* defection, which it will be if and only if they expect their partners to cooperate.

8. Slote (1989) 40.

we find one. Like local optimizing, satisficing can serve our larger plans by setting limits on how much effort we put into seeking a house at the expense of other goals that become more important at some point, given the diminishing returns of remaining on the housing market. An optimizing strategy limits how much we are willing to invest in seeking alternatives. A satisficing strategy limits how much we insist on finding before we quit that search and turn our attention to other matters.⁹

The two strategies need not be inflexible. People sometimes have reason to switch or revise strategies as new information comes in. If we seek a satisfactory house in an unfamiliar neighborhood and are shocked to find one within five minutes, we may stop the search, acknowledging the stopping rule we previously imposed on that activity. Or, we may conclude that, having formulated our aspiration level under unrealistically pessimistic assumptions, we should resume our search with a satisficing strategy revised to reflect a higher aspiration level. Or we may switch to a local optimizing strategy, spending another day or two looking at houses, then taking the best we have found so far. Or we may do both, looking until we either reach our new aspiration level or reach our time limit. In this way, the two strategies often are interactive.

Likewise, suppose we started out planning to seek the best house we could find within a one-month time limit but have so far been terribly disappointed with our options. In this case, when after two weeks we finally find a house that meets our plummeting aspiration level, we may find ourselves embracing a sadder but wiser aspiration level as a stopping rule, abandoning our original plan to seek a local optimum relative to a one-month time constraint.

Typically, the more concrete our local goals are, the more reason there is to satisfice. If we do not know exactly what we are looking for, then we usually are better off setting a time limit and then taking what we like best within that limit. But if we know exactly what we are looking for, then it is rational to stop searching as soon as we find it.¹⁰ So, having detailed information about our *goals* weighs in favor of using that information in formulating aspiration levels as stopping rules. Conversely, the more we know about our *set of alternatives*, the easier it is to identify which alternative has the highest utility, which weighs in favor of seeking local optima.

The stakes involved are also pertinent—indeed crucial. The less we care about the gap between satisfactory and optimal toothpaste, for example, the more reason we have to satisfice—to look for a satisfactory brand and stop searching when we find it. Note the alternative: instead of satisficing, we could optimize by searching among different brands of toothpaste until we find the precise point at which further search is not worth its cost.

But notice: an optimal stopping point is itself something for which we would have to search, and locating it might require information (about the probability of finding a better brand of toothpaste, for example) that is not worth gathering, given the stakes involved in the original search for toothpaste.

9. The appendix to this chapter uses graphical analysis to contrast the two kinds of self-imposed limit.

10. Jay Rosenberg tells me that before he began looking for a house, he made a list of desirable features, telling himself he would take the first house having 85 percent or more of those features. As it happens, the first house he looked at scored 85 percent. He stopped looking, bought the house, and has lived there ever since.

Against this, one might object to my assumption that we need precision in the search for an optimal stopping point. Why not seek to learn *roughly* when looking for better toothpaste is not worth the cost? In the search for a stopping point that we might graft onto the original search for toothpaste, it can be more rational to seek to be tolerably close to an optimum than to seek to be at an optimum.

But that is my point: there are cases where we do not care enough about the gap between the satisfactory and the optimal to make it rational to search for the optimal. Searching for optimal toothpaste can be a waste of time, but so can searching for the optimal moment to quit looking for toothpaste.

One way or another, satisficing enters the picture. There will be times when even the most sophisticated optimizing strategies will be inappropriate, for they require information that we may not have and that may not be worth acquiring. And a less sophisticated “all things considered” strategy will nearly always be inappropriate. Rational choice involves considering only those things that seem worthy of consideration, which is to say rational choice involves satisficing, i.e., having a stopping rule that limits how comprehensive a body of information we insist on gathering before stopping the search and turning our attention to other matters.

There is also something to be said for having a moderate disposition—a disposition that allows one to be content with merely satisfactory states of affairs. Part of the reason why is that starting a search too soon can be every bit as wasteful as stopping a search too late. Searching for a house is costly. It is costly partly because people have other goals; the time and energy you spend searching for a house could have been spent on other things. Even if you find a better house than you already have, the process of moving will also be costly. Moreover, it takes time living in and enjoying a house in order to recoup these costs. If you move every month, you will always be paying the costs and never enjoying the benefits of better housing. Moving into a house is part and parcel of a decision to stay a while, for it is only in staying that you collect on the investment of time and energy you made in moving. The general lesson is that costly transitions to preferred states of affairs require intervening periods of stability so that transition costs can be recovered and thus rationally justified. The stability of the intervening periods requires a disposition to be content for a while with what one has—to find something one likes and then stop searching.

Further, even if transition costs are relatively minor, there still can come a point when we should abandon the search for, say, a better job or a better spouse, not because such goals are unattainable or even because the transition costs are too high, but rather because such goals eventually can become inappropriate. At some point, we have to start collecting the rewards that come only when we make a genuine commitment—when we stop looking for something or someone better. We need to be able to satisfice within various local compartments (those defining our searches for spouses, jobs, and so on) in order to make our lives as a whole go well.

3. When Satisficing Is Not Rational

Slote says “choosing what is best for oneself may well be neither a necessary nor a sufficient condition of acting rationally, even in situations where only the agent’s

good is at stake.” For example, a person who is moving and must sell his house might seek “not to maximize his profit on the house, not to get the best price for it he is likely to receive within some appropriate time period, but simply to obtain what he takes to be a good or satisfactory price.” When he receives a suitable offer, he may rationally accept it immediately, even though there would be no cost or risk in waiting a few days to see if a higher offer materializes. “His early agreement may not be due to undue anxiety about the firmness of the buyer’s offer, or to a feeling that monetary transactions are unpleasant and to be got over as quickly as possible. He may simply be satisficing in the strong sense of the term. He may be moderate or modest in what he wants or needs.”¹¹

Slote does not offer an analysis of rationality. Nor do I want my argument to rest on any particular analysis of rationality. I do, however, offer this as a necessary condition of rationality: one’s choice is rational only if one does not recognize clearly better reasons for choosing any of one’s forgone alternatives. This necessary condition begs none of the questions that concern us here. It does not entail that rational choice is optimizing choice. Rather, it allows that one could rationally choose an alternative because it is satisfactory, terminating the search of one’s choice set at that point.¹² Moreover, it also allows that if one has two satisfactory alternatives, one could choose the more moderate of the two on the grounds that it satisfies a preference one happens to have for moderation.

On the other hand, although a suboptimal option may be good enough to be worthy of choice in a given case, that does not mean it is worthy of being chosen in preference to something that is clearly better. If one has two choices and one alternative is satisfactory but the other is not, then the satisfactory choice is rational because it is *better*. But suppose one has two choices and both are satisfactory. (E.g., suppose your house is for sale, and you simultaneously get two satisfactory offers, one for \$200,000 and another for \$210,000, and you prefer the larger offer.) In this case, one does not give a rationale for choosing the inferior alternative merely by pointing out that the inferior alternative is satisfactory. The inferior option is satisfactory, but since this is not a difference, it cannot make a difference either. By hypothesis, the superior option is also satisfactory.

Why, then, should we choose the superior option? Presumably because it is better. Whatever it is in virtue of which we deem that option superior is also a reason

11. Slote (1989) 1, 9, 18.

12. That is, an optimizer might choose a satisfactory option in preference to searching for better options that may never materialize. Slote, however, says we intuitively recognize the rationality of taking the first satisfactory offer even apart from real-world risks and anxieties of selling one’s house (1989, 18). However, if we are going to talk about common sense allowing a seller to immediately accept a firm offer even when the seller has the option of waiting a few days in hope of a higher offer, then we have to stick to conditions under which common sense holds sway. We do indeed have intuitions about what to do in risky situations, but we cannot, as Slote wants to do, simply *stipulate* that our intuitions regarding risky situations have nothing to do with the fact that they are risky. In the real-world housing market, to turn down an entirely satisfactory offer in quest of something better is to tempt fate. This is one reason why it is common sense, and rationally explicable common sense, for a global optimizer to be hesitant about turning down a satisfactory offer. Even from a local perspective, the expected gain from further search may not be worth risking the potential loss.

for us to choose it.¹³ Oddly, Slote denies this. It can be rational to choose the inferior option, Slote insists. Nor do we need a reason to choose the inferior option, Slote argues, because rationality does not always require people to have a reason for choosing one alternative rather than another.¹⁴ For example, Kate might rationally grab a blouse out of her closet in the morning without being able to explain why she chose that one over the similar blouses hanging beside it. To call her irrational simply because she cannot explain her choice would be a mistake.

This seems right, as far as it goes; not all choices have to be or can be explained. To deem a choice rational, however, is to imply there is an explanation of a certain kind. A person can be rational without being aware of reasons for everything she does, but the things she does for no reason are not rational, and we do not show them to be rational merely by pointing out that they were done by a rational *person*. The person who simply grabs a blouse may be choosing, perhaps rationally, to forgo the opportunity to rationally choose which of her several blouses she ends up wearing. If Kate is running late for the train, then anything that counts as a blouse will be good enough for Kate under the circumstances, so she leaves to impulse the selection from her set of blouses. (In this case, the process of searching among alternative blouses virtually vanishes—there is hardly any choice at all. If she instead gives herself a few seconds to make sure she avoids the blouses with valentine or hammer-and-sickle patterns on them, then she will be choosing within a tiny but still real local compartment.)

There may be a blouse in her closet that, given time, would emerge as best. Kate judges, however, that it is not worth her time to wait for this to happen. She is not literally compelled to simply pick something, but it serves her broader ends to forget about seeking the optimal blouse and instead just grab something out of the closet. If Kate is running late for the train, she has reason simply to grab a blouse in preference to the clearly inferior alternative of wasting precious time seeking the optimal blouse. Initially adopting an end and creating a compartment within which to pursue it is itself a goal-directed activity and, from the standpoint of the global optimizer, not one to be engaged in frivolously. Therefore, we can endorse her *method* of selecting a blouse even though we anticipate having no particular reason to endorse her actual selection.

On the other hand, *deliberately* choosing the worse over the better would be irrational, and we do not give ourselves reason to soften this verdict merely by reminding ourselves that rational people sometimes leave their choices to impulse. Rational choice theory can tell us a story about why Kate finds herself going to work in a green blouse with orange polka dots, but the story will require an implicit or explicit distinction between more and less encompassing perspectives. Without the distinction, an optimization story would be blatantly false, for she does not in fact choose the optimal blouse, and a satisficing story would have neither explanatory nor justificatory power, for the point of choosing a merely satisfactory blouse when better ones were available would remain a mystery. To see the point of what she does at the local level, we have to step back and look at her actions from a broader perspective. From a broader perspective, Kate may have good reason to simply grab a blouse out

13. Pettit (1984, 172) makes the same point.

14. Slote (1989) 21.

of the closet, knowing it will be satisfactory even if it is not her favorite. However, it cannot be rational to choose something because it is satisfactory while at the same time having a clearly better option in hand.

4. When Moderation Is Rational

When we need to decide between two satisfactory alternatives, it does not help to point out that one of them is satisfactory. We could, however, choose on the grounds that one of them is more moderate. For example, Slote says it “makes sense” for someone to desire

to be a really fine lawyer like her mother, but not desire to be as good a lawyer as she can possibly be. This limitation of ambition or aspiration may not stem from a belief that too much devotion to the law would damage other, more important parts of one’s life. In certain moderate individuals, there are limits to aspiration and desire that cannot be explained in optimizing terms.¹⁵

I agree that common sense can recognize moderate aspirations as rational, but to note this fact in an off-the-cuff way is hardly to provide an explanation of moderate aspirations. Our common-sense recognition is precisely what has to be explained. If all we have is an intuition that an act makes sense, but cannot say what the act makes sense *in terms of*, then we are jumping to conclusions if we say we were approving of the act as rational. In contrast, if we explain a show of moderation in terms of its conduciveness to overall satisfaction, then we have explained it as rational. We have not merely claimed it makes sense; rather, we have actually made sense of it. We have shown that we had reason to choose as we did, while not having better reasons to choose differently.

How, then, might we explain having moderate career goals? First, as Slote mentions, there is the issue of trade-offs. One might cultivate an ability to be content with moderate career goals not because one prefers moderate success to great success but because one cares about things other than success. Thus, one point of cultivating modest desires with respect to wealth is that it might improve a person’s ability to adhere to a satisficing strategy with respect to income, thus freeing herself to devote time to her children, her health, and so on.

There are also ways moderation can have instrumental value that do not depend on the need to make trade-offs. There can be reasons for striving to be as good a lawyer as one’s mother even if one wants to be as good a lawyer as possible. For example, a person might aim at being as good as her mother as a stepping-stone to becoming the best lawyer she can be. The modesty that enables a person to concentrate on successfully making smaller steps may eventually put her within reach of loftier goals. There is also value in concreteness. A person may have no idea how to go about becoming the best possible lawyer but may have a much clearer idea about how to become as good as her mother, because the more modest goal is more

15. Slote (1989) 2.

concrete. Further, even given two equally concrete goals, an optimizer might very well choose the lesser on the grounds that only the lesser goal is realistic. Thus, one might become a better lawyer by emulating one's highly competent mother than by wasting one's time in a fruitless attempt to emulate her superstar partner.

Finally, we can imagine moderation being a preference in itself—not just a quality of a desire but itself the thing desired.¹⁶ One might explain the cultivation of such a preference on the grounds that moderation is less distracting than extravagance, with the consequence that the moderate life is the more satisfyingly thoughtful and introspective life. In various ways, then, moderation can have instrumental or even constitutive value from the global perspective. Insofar as moderate preferences can be deliberately cultivated, their cultivation is subject to rational critique, and can thus be defended as rational.

5. When Seeking Optima Is Not Rational

To seek optima strikes us as generally a reasonable strategy, but it is not necessarily so. Section 2 noted that local optimizing can be a waste of time from a global perspective, but this is not the only circumstance that can make it inappropriate to seek optima. For one thing, a set of alternatives need not contain a well-defined optimal choice at all, let alone one that can be easily identified. To borrow a fanciful example from John Pollock, suppose you are immortal, and are also fortunate to have in your possession a bottle of EverBetter wine. This wine improves with age. In fact, it improves so steadily and so rapidly that no matter how long you wait before drinking it, you would be better off, all things considered, waiting one more day. The question is, when should you drink the wine?¹⁷

A rational *person* presumably would simply drink the wine at some point (perhaps after artificially constraining himself to drink the wine by year's end, and then picking New Year's Eve as the obvious choice within that time frame), but the person would not be able to defend any particular day as an optimal choice. Indeed, it is part of the story that no matter what day the immortal chooses, waiting one more day would have been better. There are no constraints with respect to which he can regard any particular day as the optimal choice, unless he imposes those constraints on himself.

There is something rational about choosing New Year's Eve, but the rationality lies in something other than how that day compares to the alternatives. Although the immortal could not defend choosing New Year's Eve in preference to waiting one more day, the choice is defensible in the sense that he did not have a better alternative to picking *something or other*. Indeed, picking something or other was optimal, because it was better than the only alternative, namely, sitting on the fence forever. The distinction between local and global optimizing thus allows us to explain without paradox the sense in which choosing New Year's Eve was rational. Picking something or other—and thus closing the compartment within which he seeks to set

16. I thank Mark Ravizza for this point.

17. Pollock (1984) 417.

a date for drinking the wine—was rational from the global perspective despite the fact that from within that compartment, it was not possible to have a rationale for the choice of any particular day.¹⁸

The EverBetter wine story is fantasy, of course, but it shows that we can imagine a set of alternatives with features making it inappropriate to seek the set's optimal member. Therefore, seeking optima may serve our ends, but if this is so, it is not a necessary truth but rather a contingent truth about the world and the kind of choice sets we find within it. In the EverBetter Wine case, the set of alternatives has no optimal member. Consider a more realistic story with a somewhat similar structure. Suppose a house comes up for sale in January. Out of curiosity, you take a look and find that you prefer it to the house you now own. When you look into the cost of selling your house and buying the new one, you find that the only cost you care about in the end is the cost and inconvenience of actually moving your belongings and settling into the new house. Suppose this cost, all things considered, amounts to \$1,000. Moreover, it is clear to you that such moving costs will be amply repaid over time. You can see that the stream of revenue or utility from the new house will be worth \$100 per month more than what you will receive if you stay where you are. Thus, the cost of the move will be repaid in ten months. This is hardly a wild fantasy, and so far, buying the house is intuitively reasonable. We see the point.

Now, to make the story more improbable, suppose you change houses in January and, four months later, it happens again. You find another house for sale. The move will cost another \$1,000 but the new house will be worth \$100 per month more than the one you now own. However, if you choose to move in May, that choice will make your January move retroactively suboptimal, a net loss of around \$600. Should you move?

Perhaps opportunities to move to ever better houses will surface again and again. You do not know.¹⁹ But you do know this: for any move to be optimal, something must subsequently make you stay put long enough to recover the cost of that particular move. If you keep waiting for and expecting the day when the world stops presenting you with such opportunities, and if that day never comes, then sooner or later you will have to begin turning your back on them. As in the EverBetter wine story, there is no particular point at which it is especially rational to stop moving. Indeed, whenever you finally reject an opportunity to move, it will be true that if you moved one more time, you eventually would be better off. Nevertheless, you have reason to commit yourself to being satisfied for a time with the house you have. Recall that optima are defined with respect to constraints. If you resolve in May that, once you choose, you will not look at another house for at least ten months, then choosing to move is optimal with respect to the choice set defined in part by that self-imposed constraint. Your January move will then have been a waste, but

18. Ullmann-Margalit and Morgenbesser (1977, 758–759) say one *picks* between A and B when one is indifferent between them and prefers the selection of either A or B to the selection of neither. What I call “picking” presumes the latter but not the former condition, for one could be in a picking situation even if one was not indifferent between one's alternatives. In the EverBetter wine case, one cannot find even a pair of alternatives over which one is indifferent. Even so, one still is forced to simply pick.

19. If you knew that ever better opportunities will keep coming in a steady stream, the optimal long-run strategy would be to make a big move up every ten months, skipping intervening steps.

your move in May will be worthwhile, provided that your self-imposed constraint remains firm.

Satisficing strategies strike us as reasonable in part because of contingent facts about ourselves and our world. For creatures as limited as ourselves, satisficing often makes a lot of sense. Perhaps less obvious is that the intuitive reasonableness of optimizing is no less contingent. Seizing on opportunities to make optimal moves serves a purpose partly because the real world is such that we can take for granted that there will be time between moves to enjoy our improved situation. In the real world, opportunities to improve our situation do not come along so rapidly that we find ourselves stepping higher and higher without having time to enjoy the steps along the way. The real world limits our access to opportunities to improve our situation, and if such limits did not exist, we would have to invent them. We would have to give ourselves time to enjoy our situation even if that meant rejecting opportunities to improve it.

This section explained why seeking optima is only contingently rational. The argument went beyond the idea that different local goals can come into conflict. To be sure, there can be conflicts between the pursuit of local optima and the attainment of global optima, and such occasions give us reason not to pursue local optima. This section, though, articulated a different kind of reason not to pursue local optima, because the conflict discussed in this section could occur even if one had no goals beyond, for example, living in the best possible house. The nature of the conflict is that, ironically, *seeking* the best possible house could leave us with no time to actually live in the best possible house.

6. Trade-offs among Incommensurable Values

As explained in previous sections, moderate preferences and satisficing strategies can be of instrumental value from the global perspective. Section 4 closed with the speculation that moderate preferences might even be considered essential constituents of the good life and thus have more than merely instrumental value. Satisficing strategies, however, are of instrumental value only. This is because to satisfice is to give up the possibility of attaining a preferable outcome, and giving this up has to be explained in terms of the strategic reasons one has for giving it up. Local optimizing must likewise be explained, for it, too, consists of giving something up, namely the opportunity to invest one's efforts in some other compartment.

Global optimizing, however, is not open to question and subject to trade-offs in the ways local optimizing and satisficing are. Local goals can compete with each other, but there are no goals that compete with optimizing at the global level, at least not in the arena of rationality. A global optimum is not one among several competing goals; rather, in encompassing our lives as a whole, it also encompasses our competing goals. It represents the best way to resolve the competition from the standpoint of life as a whole. Local optimizing can be a waste from the global perspective. Global optimizing cannot be.

What, then, is the nature of the global perspective? Do we ever actually assume the global viewpoint or is this merely a theoretical postulate? The answer is that we can and do assume the global viewpoint every time we do what we call "stepping

back to look at the big picture.” We do sometimes ask ourselves if the things we do to advance our careers, for example, are really worth doing. We do not spend all our waking hours looking at the big picture, of course. Nor should we, for when we look at the big picture, one thing we see is that it is possible to spend too much time looking at the big picture. Reflection is a crucial part of the good life, but it is only a part. Part of attaining a global optimum involves being able to lose ourselves for a time in our local pursuits.

Previously, we saw that, at least in imaginary cases, there can be rational choice regarding a set of alternatives even when the set has no optimal members. The lesson applies to more realistic situations as well. In particular, as Isaac Levi notes, a person torn between ideals of pacifism and patriotism need not feel that his eventual choice is best, all things considered. Rather, he may feel that his eventual choice is best according to one of his ideals and worst according to another. What we have in such a case is what Levi calls “decisionmaking under unresolved conflict of values.”²⁰ If you have several goals, none of which are subordinate to any other, and you find yourself in a situation where these goals are in conflict, the globally optimal trade-off may not exist. And such situations (involving concerns for one’s loved ones and for one’s ideals, for example) may be common.

Yet even in situations where there is no such thing as a global optimum, we can still take a global perspective. We can still look at our lives as a whole, even if nothing presents itself as optimal from that perspective. Indeed, conflict of values is precisely that from which broader perspectives emerge. We confront the big picture precisely when we stop to consider that there is more to life than pursuing a career or buying a house or raising children. It is from broader perspectives that we attempt to resolve conflicts of values, with or without an algorithm for resolving them in an optimal fashion.²¹

One might think unresolved conflict is a sign of poorly chosen values. Why should global optimizers risk adopting goals that could leave them having to make decisions under unresolved conflict? One reason is that some of our goals realize their full value in our lives only when they develop a certain autonomy, when we pursue them not as means of making our lives go well but as ends in themselves. We begin to tap the capacity of our ideals, our spouses, and our children to enrich our lives only when we acknowledge them as having value far beyond their capacity to enrich our lives. (Cherishing them becomes more than an instrumental means of making life go well; it becomes constitutive of life going well.) And goals we come to cherish as ends in themselves tend to become incommensurable.²² We may, for instance, find ourselves in a position where we cannot fight for a cause in which we deeply believe without compromising the care that our loved ones need from us and that we wholeheartedly want them to have. Nevertheless, this is the

20. Levi (1986) 13.

21. As Gibbard (1990, 321) says, we have ways of coping other than by resolving everything.

22. Seung and Bonevac (1992) distinguish *incommensurate* rankings (no alternative is best) from *indeterminate* rankings (several alternatives are tied for best). Most of what follows concerns incommensurate rankings. In contrast, most of the cases discussed in Ullmann-Margalit and Morgenbesser (1977), like the case where a shopper chooses among identical cans of tomato soup, concern indeterminate rankings.

price of the richness and complexity of a life well lived. To have both ideals and loved ones is to run the risk of having to make decisions under unresolved conflicts of value.

Because some of our values are incommensurable, we sometimes have no method by which to identify optimal trade-offs among conflicting local goals. In such cases, the goal of making life as a whole go as well as possible remains meaningful, although there may not be any course of action that unequivocally counts as pursuing it. Even if would-be global optimizers cannot identify optimal options, they can still reject alternatives that fail to further any of their goals. In particular, if no better way of resolving the conflict emerges, simply picking something or other will emerge as optimal compared to the alternative of remaining on the fence, for we eventually reject fence-sitting on the grounds that it fails to further any of our goals.

This may seem a grim picture of rational choice at the global level, but there are two points to keep in mind. First, when faced with a situation in which we must simply pick something, we are likely to have regrets about paths not taken, but we naturally adapt to the paths we take, and regret can fade as we grow into our choice. Thus, an alternative somewhat arbitrarily picked from a set within which no optimum exists can eventually come to be viewed as optimal from the perspectives of people we are yet to become, even if it could not have been considered optimal at the moment of choice. Second, this discussion of underdetermined rational choice concerns a worst-case scenario. Global optimizers carry out the highest-ranked life plan when they have one. Often, however, there is no highest-ranked plan for life as a whole and thus no well-defined global optimum; there is only a need to cope with competing and sometimes incommensurable local goals. In the worst case, no course of action unambiguously qualifies as making life as a whole go as well as possible, except insofar as it is unambiguously better to move in some direction rather than none. But this gives us enough to avoid paralysis even in the worst case. By hypothesis, simply picking something emerges as the best the agent can do, and thus to pick something is to optimize with respect to the choice of whether to spend more time sitting on the fence.

It would be natural to say rational choice is choice “all things considered.” The trouble is that we often find ourselves not knowing what to consider, and it would be bad advice to tell us to consider all things. We can consider all things within a limited range, perhaps, but the limits of that range will themselves tend to be matters of choice in large part. We start out knowing that in some sense we want each aspect of our lives to go as well as possible, yet we realize that our resources are limited and that our various pursuits must make room for each other. When looking at our lives as a whole, what is most clear is that rationally managing a whole life involves managing trade-offs among life’s various activities. If the benefits that will accrue from our various pursuits are known and commensurable, then managing the trade-offs is easy, at least theoretically; we simply maximize the sum of benefits. However, in many everyday cases, the benefits are neither known with any precision nor straightforwardly commensurable with other benefits. Even so, we can effectively manage trade-offs among particular pursuits by setting limits on how much of our lives we spend on particular pursuits. We can also set limits

on how much benefit we insist on getting from particular pursuits. To impose the latter kind of stopping rule on a particular pursuit is to embrace what I have called a satisficing strategy.

Both kinds of constraint play a role in rational choice. Why? Because if we recognized only temporal limits, say, then we would automatically spend our full allotment of time in a given compartment even when we already had an acceptable option in hand. But if we also have strategically limited aspiration within that compartment, then finding an acceptable option will trigger a second kind of stopping rule. The second stopping rule closes the compartment and diverts the unused portion of the compartment's time allotment to other compartments where our need to find an acceptable option has not yet been met. Cultivating moderate preferences may also be advantageous in a supplementary way, insofar as moderate preferences may help us adhere to the kind of limit we impose on a pursuit when we embrace a satisficing strategy.

Against the idea that our most important goals tend to become incommensurable with each other, one might suppose our global end is simply to flourish or to be happy—and that our local goals therefore *must* be commensurable in such terms. This would be a tidy climax to an otherwise rather untidy story about rational choice under unresolved conflict of values, but the tidiness would be superficial. One hardly gives people an algorithm for resolving conflicts when one advises them to be happy. What makes such advice vacuous is that flourishing and being happy cannot be concrete goals at the global level in the way that finding a house can be at the local level. Of course we *want* to flourish, but we *aim* to flourish only in an especially metaphorical sense. The fact is that we flourish not by aiming at flourishing but by successfully pursuing other things, things worth pursuing for their own sake.

Likewise, happiness can be a standard by which a life as a whole is judged, perhaps, but it cannot be a goal at which a life as a whole is aimed. We do not become happy by pursuing things there would otherwise be no point in pursuing. Rather, there must be a point in striving for a certain goal before striving for it can come to have any potential to make us happy. To aim at happiness is to aim at a state that emerges only in the course of aiming at something else.²³ So, the point about happiness and flourishing leaves us where we started, having to choose among things we value for their own sake, hoping we will be happy with our choice.

We might add that happiness derives from a variety of local sources, and the different elements of a person's happiness are not interchangeable. Our various local pleasures are not fungible; different dollar bills are all the same, functionally speaking, but different pleasures are not all the same, and are not experienced as interchangeable units of the same kind of stuff. We can find happiness in our careers and in our marriages, but the vacuum left by a shattered career cannot be filled by domestic bliss.²⁴

23. As Williams (Smart and Williams, 1973, 113) puts it, one has to want other things for there to be anywhere that happiness can come from. See also the eleventh of Butler's *Fifteen Sermons* (1874, 139).

24. I thank Nick Sturgeon for a discussion from which this point emerged. See also Stocker (1990, chap. 6) and, of course, chap. 2 of Mill's *Utilitarianism*.

7. An Infinite Regress of Perspectives?

The global perspective is the perspective encompassing our lives as a whole. Decision-making at this level disciplines the amount of time we devote to particular local compartments. It seems that we are capable of taking a perspective this broad even in worst-case scenarios where there is no well-defined global optimum. But even if we suppose we can take a perspective encompassing our whole lives, why should we suppose this is the broadest perspective we can take?

Perhaps there can be broader perspectives than what I call the global perspective. Indeed, I argue elsewhere (see chapters 3 and 4) that we do have access to a larger perspective, that there are aspects of morality that we cannot appreciate except from a larger perspective, and that it can be rational to try to achieve this perspective. Yet, it would be unrealistic to suppose there is an infinite regress of levels. We need not prove that an infinite regress is impossible, but insofar as an infinite regress is unrealistic, it is important to show that my theory does not *presuppose* an infinite regress.

The threat of infinite regress arises in the following way. I said we cannot spend all our time looking at life as a whole; we must be able to lose ourselves (or perhaps I should say, find ourselves) in our local pursuits. How much time, then, should we spend pondering conflicting values? How much time should we spend looking at life as a whole? From what perspective do we choose to limit the amount of time we spend looking at our lives from the global perspective? Perhaps we need a “superglobal” perspective in order to answer these questions. After all, how could we decide how much time to spend at a given level unless we did so from a still more encompassing perspective?²⁵

It seems my theory can explain the time we allot to a given perspective only by supposing that we retreat to a broader one, ad infinitum. Not so. The theory presumes no such retreat. There are simpler, more realistic ways to explain the amount of time we spend looking at life as a whole.²⁶ Here are two.

1. There are things, like sleeping, that we do as the need arises. Since we do not *decide* how much time to spend sleeping, we do not decide from a broader perspective, either. Indeed, we might be better off sleeping as we feel the need rather than trying to set aside a calculated amount of time for sleep. Perhaps the same holds true of the activity of looking at life as a whole. Insofar as our purpose in looking at life as a whole is to resolve conflicts arising between various aspects of our lives, so that life as a whole may go well, there will come a time when taking a global perspective has served its purpose. At that time, the compartment in our lives reserved for the activity of resolving local conflicts naturally closes until subsequent conflict forces it open again. There is no residual conflict awaiting resolution at a higher level.

25. Holly Smith (1991) worries about the same sort of problem.

26. The simplest way to explain time spent at the global level would be to say we take whatever time we need to consider *everything*. The trouble is that we do not have time to consider everything that may be relevant to life as a whole, any more than we have time to consider everything that may be relevant to purchasing a house. The explanation will have to be more complicated than this; hence the line of thought pursued in the following text.

Thus, the question of how much time to spend in contemplation need not itself require contemplation. Rather, we take whatever time it takes to genuinely resolve a conflict, or else we reach a point where we must simply pick something. More generally, we stop contemplating when we judge that pursuing our local goals has come into conflict with—and has become more important than—the activity of thinking about how to juggle them. (For example, we would not dwell on the big picture if we were starving. Conflicts are rarely so important that contemplating them could preempt securing our immediate survival.) In this scenario, we are driven *to* the global level by local conflict and eventually are driven *from* that level by a need to get on with our lives.²⁷ The question of how much time to spend looking at life as a whole resolves itself.

2. We also can imagine a scenario in which the question does not resolve itself but is instead answered by deliberate calculation, in the same way we could imagine deliberately calculating how much time to spend sleeping. Could we make a conscious decision of this kind without taking a superglobal perspective? Yes, we could. Consider that contemplation is an activity that must find its place in our lives along with other activities. For example, I might spend the month of July in a rented cabin, not doing anything to pursue my career, but just thinking about why I ever wanted to be a philosopher and about whether my original reasons still hold. This compartment in my life is reserved for contemplating my career. It is separate from the compartment or compartments within which I actually pursue my career. I also have a compartment, similar in many respects, within which I contemplate life as a whole. But although the subject I contemplate is the whole of my life, the contemplation itself is not. The contemplation is only one of many activities about which I care.

Now, if I need to decide how much time to reserve for contemplating life as a whole, I take a global perspective, trying to gauge how important such contemplation is to my life as a whole. Notice, then, what is unique about the compartment I reserve for the activity of contemplating my life as a whole. The compartment is unique because its boundaries are set by the activity that takes place within it. In the course of contemplating life from the global perspective, I decide how much time to reserve for any given activity, including contemplation in general and contemplating life as a whole in particular. In this second scenario, as in the first, there is no residue. Once we finish making the kinds of decisions we make at the global level, we are done. No boundary-setting issue is left to await resolution at a higher level.

I have outlined two possibilities. In the first case, we use whatever time we need for resolving conflicts, subject to preemption by activities that in the short run are more important than conflict resolution. In this case, no decision is required.

27. We also can be driven to a global perspective by the resolution of conflicts. Thus, when we finish a major project that has forced other pursuits to take a back seat, we often take time to evaluate self-imposed constraints and decide how to divide our extra time among previously neglected projects. And what drives us from the global perspective is the eventual resolution of a local conflict between savoring the big picture (a satisfying activity indeed when just finishing a major project) and the need, say, to start making dinner.

The discipline is automatic. In the second case, we discipline the compartment from within, as our contemplation of trade-offs leads us to conclude that we should reserve time for contemplation along with our other local activities. Therefore, we do not need a superglobal perspective to decide how much time to reserve for the activity of taking a global perspective. Such decisions are precisely the kind we make from the global perspective itself, when we need to make them at all.

8. Conclusion

This essay sets out part of a normative ideal of rational choice suitable for the kind of beings we happen to be, beings who would only hurt themselves if they tried to maximize their overall utility in every waking activity. This essay defines satisficing and local optimizing as strategies for pursuing goals within constraints that are in part self-imposed. Satisficing emerges not as an alternative to optimizing as a model of rationality, but rather as an alternative to local optimizing as a strategy for pursuing global optima.

Under normal conditions, we employ a combination of heuristics, such as (1) compartmentalizing our pursuits so as to narrow the scope of any particular optimization problem to the point where our limited knowledge becomes sufficient to identify an optimal solution; (2) accepting self-imposed constraints for the same reason as well as to keep particular pursuits from preempting more important ones; and (3) satisficing, which has the effect of closing compartments as soon as they serve the purpose for which they were created. Under normal conditions, where we lack the information we need to assign probabilities and utilities, this combination of strategies is more effective at making our lives as a whole go well than the alternative of plugging guesswork into a formula for maximizing expected utility. Thus, it is no wonder we so rarely find ourselves trying to calculate expected utilities, for the truth is that we usually have better things to do.

When goals are in conflict, there may not be any well-defined sense in which one way of resolving the conflict is, from the viewpoint of one's life as a whole, better than the alternatives. Of course, we do well to cultivate moderate preferences, so as to reduce the frequency and severity of conflicts of value. But at the same time, there are limits to what one should do to avoid situations of underdetermined choice, for the risk of finding oneself in such situations is a risk one assumes in the process of becoming rationally committed to particular ends as ends in themselves. A life with no regrets (about decisions made under unresolved conflict) is preferable, all other things equal, but if the lack of regret is purchased at a cost of not having goals that can come into irresolvable conflict, the price is too high. A person who adopts a number of goals as ends in themselves risks finding himself in situations where global optima do not exist, but there are reasons why a global optimizer would take that risk.

Admittedly, these conclusions about rational strategy are not neat and tidy, certainly not in comparison to the simple maximization model. But tidying up the conclusions at the expense of realism would be a mistake, for the conclusions are meant to be about us, not about mathematically tractable caricatures of us. Rational choice

theory developed along the lines indicated here has more power than simple maximization models to explain why we live as we do, but it does not thereby become merely a self-congratulating description of how we live. Rather, it sets out a normative ideal of rational choice that it would be natural and healthy for us to try to live up to.

Appendix: The Difference between Satisficing and Local Optimizing

I distinguished between satisficing and local optimizing as stopping rules. Some readers may find it helpful to consider a graphical representation of that distinction. We might represent a choice among alternatives in two-dimensional Cartesian space with Utility on the y-axis and our set of alternatives arrayed along the x-axis. If we know the shape of the utility curve, we simply pick the highest point. No controversy arises. (See fig. 1.1.)

This essay concerns what to do when we are looking at a blank; that is, we may suppose there is some curve or other, but often we do not know what it looks like. (See fig. 1.2.)

Further, suppose we look at our lives from a global perspective, wanting life as a whole to go well. What do we see? We don't see one big graph, blank or otherwise. Rather, we see a bunch of little graphs, some of which are more or less blank. The question then arises: within a particular compartment, how do we make decisions when we do not know the utility function's shape? The answer is that we search the set of alternatives. We see how much utility a_1 has. We see how much a_2 has, and so on. And since other decisions (searches) are also calling for our limited resources, we pick something at some point.

At what point do we rationally stop searching and pick something? The answer is, we impose two kinds of constraints on our search of the particular local utility space. We impose vertical constraints on how many alternatives we will consider (or if we defined the x-axis differently, on how much time or money or other resources we invest in the search). In other words, we operate with limited *inputs*. Or we impose horizontal constraints on how much utility we insist on getting before we stop searching. In other words, we operate with limited aspirations, limits on aimed-at *output*. Or we do both. Then, when we run up against either kind of limit, we stop searching in that local utility space, pick something, and turn our attention to some other local utility space. (See fig. 2.3.)




FIGURE 2.1. Searching among alternatives with known utilities.




FIGURE 2.2. Searching among alternatives with unknown utilities.




FIGURE 2.3. Two stopping rules contrasted.

In figure 2.3, the horizontal line represents the point at which $U=U^*$, where U^* is the level of U with which the agent will be satisfied. As mentioned earlier, we can let the expected utility of continued search equal the probability of finding a better option, $P(fbo)$, multiplied by the utility of finding a better option, $U(fbo)$, minus the cost of further search, $C(fs)$. In that case, the vertical line in figure 2.3 represents the point along the x-axis at which, the agent judges, it becomes true that $P(fbo)U(fbo) - C(fs) \leq 0$, where $P(fbo)U(fbo)$ is the expected utility of further search and $C(fs)$ is the cost of further search.

In figure 2.3, vertical constraints are constraints on inputs, and define the search as local optimizing, taking the best alternative we discover prior to hitting that constraint. Horizontal constraints are constraints on aimed-at output, and define the search as satisficing, taking the first alternative we find with that high a utility. The two strategies can be employed simultaneously, of course. I argued that this is how we actually live and that a global optimizer would have no reason for wanting to do things differently.

The vertical constraints partition our various activities in terms of how much of our total resources are allotted to those activities. Note that we do not need to be able to prioritize our activities in order to ration our resources among them. If necessary, we can arbitrarily set vertical constraints on how much (time, money, and so on) we are willing to spend within a particular compartment. As illustrated by the story about imposing time constraints on our search for optimal toothpaste, we tend to be satisficers when gathering information about where to set vertical constraints. The



less we know about what our different endeavors mean to us, relatively speaking, the more arbitrariness there will be when we set the vertical constraints that delimit the different compartments. By the same token, the more comparability we have in terms of the relative importance of our different activities, the less arbitrary will be the boundaries we draw between them.

This chapter revises “Rationality within Reason,” *Journal of Philosophy* 89 (1992): 445–66. Reprinted by permission of the Journal.

