

---

Rationality within Reason

Author(s): David Schmidtz

Source: *The Journal of Philosophy*, Sep., 1992, Vol. 89, No. 9 (Sep., 1992), pp. 445-466

Published by: Journal of Philosophy, Inc.

Stable URL: <https://www.jstor.org/stable/2941111>

---

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact [support@jstor.org](mailto:support@jstor.org).

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



is collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Philosophy*

JSTOR

---

---

# THE JOURNAL OF PHILOSOPHY

VOLUME LXXXIX, NO. 9, SEPTEMBER 1992

---

---

## RATIONALITY WITHIN REASON\*

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 which needs me and which 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 the present case, however, I can only guess at the possible outcomes of either course of action. Nor do I know their probabilities, or 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. Explaining what they are is the purpose of this paper.<sup>1</sup>

\* I thank Simon Blackburn, Jim Child, Jules Coleman, Tyler Cowen, Peter Danielson, Walter Glannon, Mark Kingwell, Ruth Marcus, Mark Migotti, Christopher Morris, Jan Narveson, Alan Nelson, Philip Pettit, Ken Presting, Mark Ravizza, Mike Resnik, Lainie and John Ross, Stephen Scott, Michael Slote, Elizabeth Willott, and audiences at Cornell and the University of North Carolina/Chapel Hill for especially helpful comments.

<sup>1</sup> Further questions concern the relationship between rational choice and moral agency. Do we have reasons to be moral? If so, how do they relate to our reasons

## I. RATIONAL CHOICE

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 a search for alternatives upon concluding that one has identified the best available alternative. Although optimizing involves selecting what one judges is best, however, it need not involve judging what is best "all things considered," because sophisticated optimizers recognize that considering all alternatives is not always worth the cost. There may be constraints (temporal, financial, and so on) on how much searching they 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.<sup>2</sup>

Satisficing, in contrast, involves terminating the search for alter-

---

to be rational? Although I will not address these questions in this paper, they are ultimately the questions which concern me most and which motivate the larger project of which this paper is a part.

<sup>2</sup> Thus, if options emerge serially, one could choose an option in preference to the alternative of 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,  $P(fb)$ , multiplied by the utility of finding a better option,  $U(fb)$ , 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  $P(fb)U(fb) - C(fs)$  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.

natives 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*.<sup>3</sup> Unlike the optimizer, who stops searching when she has either considered all her options or has run up against things like time constraints, the satisficer stops the search upon identifying an alternative as good enough.<sup>4</sup> For example, suppose you enter a cafeteria seeking a nutritionally balanced and reasonably tasty meal. You walk 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 that 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. So, neither rule is reducible to the other. (One could employ both stopping rules simultaneously, of course, resolving to stop as soon as one finds a satisfactory alternative or runs out of time or has considered all available alternatives—whichever comes first.) Nor 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 it is satisfactory, not because they calculate that stopping the search at that point would maximize utility.

<sup>3</sup> I borrowed the term 'satisficing' from Michael Slote, *Beyond Optimizing: A Study of Rational Choice* (Cambridge: Harvard, 1989), p. 5, but my characterization of satisficing differs from his. It turns out that my characterization of satisficing as a stopping rule was anticipated by Herbert Simon in *Models of Thought* (New Haven: Yale, 1979), p. 3. In the same passage, Simon also anticipates one of my main complaints about maximization strategies, namely, that they require information that decision makers often do not have.

<sup>4</sup> I do not know of any precise way to characterize 'good enough'. Options that promise disease, imprisonment, or premature death are typically held in low esteem, however, so the notion has certain objective elements. But what people consider good enough also 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, for 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 one deems one's 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.

With this characterization of satisficing in mind, we can now clarify the difference between satisficing and *moderation*. Consider that being moderate is contrasted, not with optimizing, but with being immoderate. The distinction between being moderate and being immoderate is a distinction concerning how much one wants. It is a distinction between different ends rather than between different ways of pursuing them. Being a satisficer does not entail that the bundle of goods one would deem satisfactory is of moderate size. 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 by the age of thirty.) Likewise, that one is an optimizer does not entail that one is immoderate. A person could easily be both a moderate and an optimizer, for the bundle of goods required to maximize her satisfaction may well be of moderate size.

#### II. WHEN SATISFICING IS RATIONAL

There is an apparent incongruence between the theory and practice of rational choice. Theory models rational choice as optimizing (sometimes even “all things considered”) choice, yet in practice satisficing is ubiquitous. We could explain the incongruence away by saying that when people think they are looking for something satisfactory, what they are really looking for is something optimal. But I think 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 that we can recognize in ourselves. 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. While a given pair of goals might have little to do with each other, it is also true that a given goal might be encompassed by another. In other words, the point of the narrower goal—the reason for its being a goal—is that it is part of what one does in pursuit of a broader goal. For example, I might want to upgrade my wardrobe because I care about my appearance because I want a promotion because I care about my career. Suppose I find that the multiplicity of goals that I pursue are ultimately encompassed and rationalized by a goal of making my life as a whole go well. To mark the difference in breadth between my

concern for my life as a whole and my concern for particular aspects of my life (such as my health or my career), let us say one seeks a *local* optimum when one seeks to make a certain aspect of one's life go as well as possible. One seeks a *global* optimum when one seeks to make one's life as a whole go as well as possible.<sup>5</sup>

Note that 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 often make decisions under conditions where the set of constraints imposed on us by external circumstances is not determinate enough to yield a well-defined optimum. When this happens, we often impose further constraints on ourselves.

For example, in fleshing out the problem of locating a house, we have 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 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.<sup>6</sup>

In other words, if we look at life as a whole, we can see that life as a whole will go better if we spend most of it pursuing goals that are

<sup>5</sup> I borrowed the terms 'local' and 'global' from Jon Elster, although when rereading his text I noticed that the way he uses the terms bears little resemblance to the way they are used here. He says the definitive difference between locally and globally maximizing machines is that the latter, unlike the former, are capable of *waiting* and *indirect* strategies. See Elster, *Ulysses and the Sirens: Studies in Rationality and Irrationality* (New York: Cambridge, rev. 1984), p. 9.

<sup>6</sup> Jules Coleman has pointed out to me that what David Gauthier (*Morals By Agreement* (New York: Oxford, 1986), p. 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.

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 is also 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, we pursue narrower goals 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. (External constraints, of course, are salient even at the global level.) Self-imposed constraints can be applied only *to* narrower pursuits and can be applied only *from* the perspective of a more encompassing pursuit. Our broader objectives constrain what we are willing to spend on upgrading a wardrobe. They constrain what we will do for the sake of our appearance, or for the sake of a promotion, or for the sake of a career.

Armed with the distinction 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 that “[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 . . .” (*op. cit.*, p. 40). 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” (*op. cit.*, p. 40). 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. Thus, satisficing is a big part of a global optimizer’s daily routine. 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 certain 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. Looking for a house competes with our other goals for our time and energy. Or we might look for a satisfactory house and cease looking when we find one. Most of the

people I have asked say they would optimize within constraints, but would not deem satisficing irrational. Like local optimizing, it 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 places limits on how much we are willing to invest in seeking alternatives. A satisficing strategy places limits on how much we insist on finding before we quit that search and turn our attention to other matters.

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

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

I suspect that 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 something that fits the bill.<sup>7</sup> 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

<sup>7</sup> 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 that had 85% or more of those features. As it happens, the first house he looked at scored 85%. He stopped looking, bought the house, and has lived there ever since.



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 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. 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 one does 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 whose acquisition may not be worth the trouble. And a less sophisticated “all things considered” strategy will nearly always be inappropriate. Rational choice involves considering only those things which seem worthy of consideration, which is to say it 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. Consider 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. Local optimization processes have their own ends, but 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. We need to be able to satisfice within various local compartments (those defining our searches for spouses, jobs, and so on) in order to be able to make our lives as a whole go as well as possible.

### III. WHEN SATISFICING IS NOT RATIONAL

Slote says that “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” (*op. cit.*, p. 1). 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” (*op. cit.*, p. 9). Once the seller 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” (*op. cit.*, p. 18).

Slote does not offer an analysis of rationality. Nor shall I. I do, however, offer this as a necessary condition of rationality: one’s choice is rational only if one did not recognize clearly better reasons

for choosing any of one's foregone alternatives. This condition does not beg any 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.<sup>8</sup> 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 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 cannot rationalize 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 for us to choose it.<sup>9</sup> 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

<sup>8</sup> That is, an optimizer might choose a satisfactory option in preference to searching for better options that might never materialize. Slote, however, says we intuitively recognize the rationality of taking the first satisfactory offer even in abstraction from the real-world risks and anxieties of having to sell one's house (*op. cit.*, p. 18). But my view is that, if we are going to talk about common sense allowing the seller to accept the firm offer immediately even though the seller has the option of waiting a few days in hope of a higher offer, then we have to stick to the conditions under which common sense holds sway. We do indeed have intuitions about what we should 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 in the real world such situations are risky. In the real-world housing market, to turn down an entirely satisfactory offer in quest of something better is to court disaster, to tempt fate, as it were. 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.

<sup>9</sup> Philip Pettit makes the same point in "Satisficing Consequentialism," *Proceedings of the Aristotelian Society*, supplement to LVIII (1984): 165–76, here p. 172.

reason for choosing one alternative rather than another (*op. cit.*, p. 21). For example, a rational person might 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.

I agree with this, 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 choose rationally which of her several blouses she ends up wearing. If she is running late for the train, then under the circumstances anything that counts as a blouse will also count as satisfactory, 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 blouses with valentines or political slogans on them, then she will be choosing within a small but real local compartment.)

Thus, it can be rational to forgo the opportunity to choose rationally, because searching for reasons to choose one alternative over another might be a waste of time from a broader perspective. Nevertheless, it definitely is a condition of rational *choice* that the choice in question is made for a reason. Deliberately choosing the worse over the better is irrational, and this is not mitigated by the fact that rational people sometimes leave their choices to impulse.

Rational-choice theory can tell us a story about why the commuter 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 be able to see the point of what she does at the local level, we have to be able to step back and look at her actions from a broader perspective. From a broader perspective, it may be perfectly reasonable for her to cut short a particular search because she was looking for something satisfactory and found it, but

it cannot be rational to choose something because it is satisfactory while at the same time recognizing a clearly better alternative.

#### IV. WHEN MODERATION IS RATIONAL

As I said earlier, we cannot choose between two satisfactory alternatives on the grounds that one of them is satisfactory. We could, however, choose on the grounds that one of them is more moderate. Consider an example of Slote's. He 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 . . ." (*op. cit.*, p. 2).

Common sense indeed recognizes that moderate aspirations can be 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 would be jumping to conclusions if we said 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 that 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, there is the issue of trade-offs mentioned by Slote. 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 in which 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 something more lofty. 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 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 at least conceive of moderation being a preference in itself—not just a quality of a desire but itself the thing desired.<sup>10</sup> 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.

#### V. GLOBAL PERSPECTIVE AND UNDERDETERMINED RATIONAL CHOICE

Like moderate preferences, satisficing strategies can be of instrumental value from the global perspective. Satisficing strategies, however, can be 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 that 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 of time from the global perspective, but global optimizing cannot.

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 view-

<sup>10</sup> I thank Mark Ravizza for this point.

point 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 we can 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.

To speak of attaining a global optimum, however, and to endorse global optimizing as a normative ideal, is not to presume that there is an algorithm for identifying global optima. In fact, not every set of alternatives even contains a well-defined optimal choice, let alone one that can be easily identified. To borrow a fanciful example from John Pollock,<sup>11</sup> 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?

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 built into the example 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 about choosing New Year’s Eve that is rational, but it consists of 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 the optimal selection from the set of options whose only other member consisted of 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

<sup>11</sup> “How Do You Maximize Expectation Value?” *Nous*, xvii (1984): 409–21, here pp. 417ff.

rational. Picking something or other—and thus closing the compartment within which he seeks to set 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 rationalize the choice of any particular day.<sup>12</sup>

This lesson also applies to more realistic situations. In particular, as Isaac Levi<sup>13</sup> 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 “decision making under unresolved conflict of values” (*ibid.*, pp. 13ff). If you have several goals, none of which is 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 rather common.

Yet, even in situations like those envisioned by Levi, where there is no such thing as a global optimum, we still can take a global perspective. We still can 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.<sup>14</sup>

One might think unresolved conflict is a sign of poorly chosen values. Why should would-be 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

<sup>12</sup> Edna Ullmann-Margalit and Sidney Morgenbesser, in “Picking and Choosing,” *Social Research*, XLIV (1977): 757–85, here pp. 758–9, 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 something or other” presumes the latter but not the former condition, for I think one could be in a picking situation even if one was not indifferent between one’s alternatives. This is most clearly illustrated by the “EverBetter Wine” case, where one has an infinite number of alternatives, each better than the previous one. In this case, one could not find even a pair of alternatives over which one was indifferent, but one is nevertheless forced to simply pick.

<sup>13</sup> *Hard Choices: Decision Making Under Unresolved Conflict* (New York: Cambridge, 1986).

<sup>14</sup> As Allan Gibbard once said, we have ways of coping other than by resolving everything. See *Wise Choices, Apt Feelings* (Cambridge: Harvard, 1990), p. 321.



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 that we come to cherish as ends in themselves inherently tend to become *incommensurable with each other*. 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 which our loved ones need from us and which we wholeheartedly want them to have. Nevertheless, this is the 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 be 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 net benefits. In many of the everyday cases I have been discussing, however, the benefits are neither known nor commensurable with other benefits. Even so, we can effectively manage trade-offs among our particular pursuits by setting limits on how much of our lives we will 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 had already found an acceptable option. 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, I think, 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 things other than flourishing, 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 property that can emerge only in the course of aiming at something else.<sup>15</sup> 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.

I might add that happiness derives from a variety of local sources, and the different elements of a person's happiness are not interchangeable. We can find happiness in our careers or in our marriages, but the vacuum left by a shattered career cannot be filled by domestic bliss. Our various local pleasures are not fungible. They do not admit of an overall measure suitable for use at the global level.<sup>16</sup>

#### VI. THE BROADEST PERSPECTIVE

The global perspective is the perspective that encompasses our lives as a whole. Decision making at this level disciplines the amount of time we devote to particular local compartments. I think 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 that we can take a perspective encompassing our whole lives, why should we suppose that this is the broadest perspective we can take?

<sup>15</sup> As Bernard Williams puts it, one has to want other things for there to be anywhere that happiness can come from. See "A Critique of Utilitarianism," in J. J. C. Smart and Bernard Williams, *Utilitarianism: For and Against* (New York: Cambridge, 1973), p. 113.

<sup>16</sup> I thank Nick Sturgeon for a discussion from which this point emerged.

Perhaps there can be broader perspectives than what I have called the global perspective. Indeed, in the larger project of which this paper is a part, I hope to show that we do have access to a larger perspective, that there are some aspects of morality that we cannot appreciate except from this larger perspective, and that it can be rational to try to achieve this perspective. On the other hand, I think it would be highly unrealistic to suppose there is an infinite regress of levels. There is no need to prove that an infinite regress is impossible, but I would like to be able 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 decide what constraints to apply to the amount of time we spend looking at our lives from the global perspective? Perhaps we need a “super-global” 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? It seems that my theory can explain the time we allot to a given perspective only by supposing that we retreat to a broader one, *ad infinitum*.

But 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.<sup>17</sup> First, 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 as well as possible, there will come a time when taking a global perspective has served its purpose. At that time, the compartment reserved for the activity of resolving local conflicts naturally closes until subsequent conflict forces it open again. There is no

<sup>17</sup> The simplest way to explain the amount of time spent at the global level would be to say that we take whatever time we need to consider *everything*. The trouble is that we do not have time to consider everything that might be relevant to life as a whole, any more than we have time to consider everything that might be relevant to the purchasing of a house. The explanation will have to be more complicated than this; hence the line of thought pursued in the following text.

residual conflict awaiting resolution at a higher level. Thus, the question of how much time to spend in contemplation need not itself require contemplation. Rather, we take whatever time it takes to 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 pre-empt 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.<sup>18</sup>

But we can also imagine a second kind of scenario in which the amount of time we spend looking at life as a whole is determined by deliberate calculation, in the same way that we could imagine deliberately calculating how much time to spend sleeping. Could we make a conscious decision of this kind without taking a super-global 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 June 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 that activity 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

<sup>18</sup> We also can be driven to a global perspective by the resolution of conflicts. Thus, when we finish a major project that had 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.

decide how much time to reserve for any given activity, including contemplation in general and contemplating life as a whole in particular. In this scenario, like the previous one, no boundary-setting issue is left to await resolution at a higher level.

I have outlined two possibilities here. In one case, I use whatever time it takes to resolve conflicts, subject to pre-emption by activities that in the short run are more important than conflict resolution. In this case, no decision is required. The discipline is automatic. In the second case, I discipline the compartment from within, as the process of contemplating trade-offs culminates in time being reserved for contemplation along with other local activities. Therefore, I do not need a super-global perspective to decide how much time to reserve for the activity of taking the global perspective. Such decisions are precisely the kind I make from the global perspective itself, if I need to make them at all. Unless we introduce something that competes with the goal of making life as a whole go as well as possible (such as, perhaps, the recognition of moral obligations), there is no reason to step back from a global perspective to something even broader.

#### VII. CONCLUSIONS

This paper sets out a normative ideal of rational choice that is suitable for the kind of beings we happen to be, beings who would only hurt ourselves if we tried to maximize our overall utility in every waking activity. It 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 pre-empting 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 make any attempt 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 better than the alternatives from the viewpoint of one's life as a whole. 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 under-determined choice, for the risk of finding oneself in such situations is a risk we assume in the process of becoming rationally committed to particular ends as ends in themselves. A life without regrets (i.e., 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 unresolvable conflict, the price is too high. To adopt a number of goals as ends in themselves is to risk finding oneself in situations where global optima do not exist, but there are reasons why a global optimizer would take that risk.

I admit that the theory is not particularly neat and tidy, certainly not in comparison to the standard model. But tidying up the theory at the expense of realism would be a mistake, for the theory is meant to be about us, not about mathematically tractable caricatures of us. I believe rational-choice theory developed along the lines indicated here has more power than standard maximization models to explain the ways in which we actually live, but it does not thereby become merely a self-congratulating description of how we live. Rather, it remains (or becomes) a tool for evaluating and criticizing the ways in which we actually live. It sets out a normative ideal of rational choice that it would be natural and healthy for us to try to live up to.

DAVID SCHMIDTZ

Yale University