American Philosophical Quarterly
Volume 46, Number 4, October 2009

RESTING CONTENT: SENSIBLE SATISFICING?

Patricia Greenspan

Suppose I am now making plans for next summer's vacation. I can spend a week in Rome or on the Riviera, but not both. Either choice would be excellent, but after weighing various pros and cons, I decide that for my purposes Rome would be significantly better. If I am rational, then, it seems I must choose Rome.

It is an assumption of standard decision theory that rationality requires maximizing: trying to get the maximum amount of whatever form of value we are after (usually construed as "utility"). An alternative has been proposed, under the heading of "satisficing"—being satisfied with what suffices, as it were: settling for an option that is "good enough." But this may seem rational only when there are costs to determining which option is best that diminish the value of choosing it, to the point where the choice of a merely satisfactory option really amounts to maximizing. Where there is no serious cost, even in time and effort, to getting hold of something better—a vacation in Rome rather than the Riviera—how could it be rational to turn it down?

If we pay attention to the temporal standpoint from which a choice is made, though, satisficing makes good sense. It accords with our common appeal to thresholds: adequate levels of satisfaction or value, such that getting above them is not necessary, though it might be nice. Once we have reached a threshold, it is rational—meaning rationally permissible—to stop. Pushing further toward the best may also be permissible but is not rationally required, if we already have a good enough option in hand. So if offered a chance to move to Rome while already settled happily on the Riviera, I would not be irrational to turn the offer down, even if I grant that accepting it would make my vacation even better.

Thresholds support a notion of *optional* reasons—all-things-considered reasons that exhibit something valuable about an option but do not require it. If I have already met a value threshold, there is no need for me to act on a reason to exceed it that does not offer any criticism of the status quo (see Greenspan 2005 and 2007). In other words, sometimes a reason, rather than requiring an option (even prima facie), just presents it as rational to choose—though one might also rationally turn it down, even in the absence of reasons against it. "Rational," once again, is interpreted here as "rationally permissible"—"within reason," or "not *ir*rational," rather than uniquely or ideally rational.

The aim of this paper is to defend satisficing in commonsense terms—to explain why it should count as rational—in cases where a threshold has been met. These might be thought of as "bird-in-hand" cases, in con-

trast to cases of choice among several options at roughly equal or remote distance in time. What is involved in bird-in-hand cases is popularly referred to as "standing pat": remaining at a level already reached rather than attempting to go higher.

The paper attempts just to argue that satisficing makes sense in bird-in-hand cases, that there is a coherent rationale for it, not to refute maximizing theories. The point is not to represent satisficing as a possible basis for an overall conception of rational choice that might compete with standard decision theory, but just to explain why it amounts to a reasonable way of making a choice in common cases—in principle, that is, not merely because we need to take shortcuts or employ some sort of heuristic in order to save time or conserve energy or make do with limited cognitive resources. But the defense of satisficing will turn on limitations of another sort—in our control over subjective factors in decision making that affect the value of our options.

The notion of satisficing originally surfaced in economics in connection with constraints on human informational capacities (see Simon 1955). In recent psychology it is treated as one among several methods we have for "fast and frugal" decision making that approximate to an optimal decision (see, e.g., Gigerenzer and Selten 2002). However, some recent authors in moral philosophy have defended "in principle" variants of satisficing that are not based on human cognitive limitations but on parallels to ethical theory.

One line of thought comes from consequentialist ethics, where insistence on maximizing yields a version of utilitarianism that intuitively seems too demanding, with a requirement that one do what has the best consequences for everyone, even if relieving others' suffering worldwide would entail dropping to subsistence level oneself. An alternative that requires only meeting a threshold of "good enough" consequences has been applied to ethics and then read back onto individual rationality; but in a case where one is already well above the relevant value-threshold, this would seem to permit making things worse, which intuitively seems irrational (cf. Slote 1985 and Bradley 2006). In contrast, an account based on bird-in-hand cases would take the threshold to be set no lower than the level already reached.

Another version of rational satisficing has been proposed on the basis of virtue ethics, as entailed by the virtue of moderation (see Slote 1989). But this is extreme in its own right, since it would *prohibit* maximizing. By contrast, the proposed account of bird-in-hand satisficing would simply allow satisficing as a rationally permissible alternative to maximizing.

The defense of satisficing in this paper will not work from parallels in ethical theory or from human cognitive limitations. It will later make reference to some of each, but its basis is in concerns of moral psychology with a bearing on practical rationality. For even without limitations in our ability to solve complex choice tasks, piling on such tasks in order to maximize would tend to undermine our capacity to enjoy the results of past choices. Often we are well advised to rest content with what we have rather than subject it to further comparison shopping.

A proponent of maximizing might be tempted just to add enjoyment or some other subjective state such as contentment to the mix of values making an option ideal, but there are more plausible alternatives, such as avoiding a kind of waste, that fit the satisficing model. What is in question here, more generally, are habits or patterns of choice that also include requisites of rational planning agency. We place a value on reliability, or consistency over time, that can pull against too much comparison of further options. But we need to balance this against competing considerations such as spontaneity or flexibility, and at this point the burden shifts:

maximizing begins to look like the position that needs defense. Requiring an optimal mix of such factors would be unrealistic, partly in light of our limited control over attitudes such as contentment that serve to define both the ideal and the threshold of adequacy.

This general line of defense for resting content with less than the best, to be explained and developed further in what follows, need not be tied to the technical term "satisficing," though it is worth noting that the term was originally meant to apply to instances of serial choice, in which options are encountered in succession. At a certain point in time these constitute bird-in-hand cases. But the aim of the paper is essentially just to defend everyday threshold-reasoning as more than a mere labor-saving device in deliberation. Its explanation in terms of multiple competing general values, and ultimately in terms of the value of attitudes like contentment, departs from accounts of satisficing in the literature. Later in the paper some alternatives will be considered and shown to be either inadequate to handle the relevant cases or compatible with the account proposed. Bird-in-hand satisficing might also just be dismissed as an instance of a misleading heuristic, assigning more weight to losses than to gains foregone, but by the end of the paper we should be able to see why this form of satisficing makes sense in rational terms.

I. When Maximizing Is Required

Let us begin by looking again at the case of choice between vacation options, in order to isolate a version of the case in which satisficing would indeed seem rationally unintelligible. I am debating among options for a trip: to tour Rome in the summer heat or instead to relax on the Riviera. I acknowledge that my reasons to go to Rome are stronger than the reasons I have to spend time on the Riviera: I can always relax at home, whereas I never have managed to see the Colosseum and other ancient Roman sites. Could it be rational, in light of that, for me to choose the Riviera?

My options here are assumed to be commensurable—we can reach an overall judgment of their comparative value (or their likely value to me this summer)—and we can also think of my reasons for each of them as comparable in terms of strength. But the presence of varied specific grounds for choice in complex cases like this one can still support an alternative to maximizing overall value. I might acknowledge that Rome would be better than the Riviera overall, and that my reasons for going there are stronger, but decide to go to the Riviera instead, because spending a week there would be more relaxing. A decision on my part to prioritize a particular value (or form or aspect or dimension of value) would let me *discount* considerations of overall value and the reasons that reflect them. I have a higher-order option regarding my basis for making the choice: I can legitimately decide to stress relaxation over cultural enrichment, say, while still taking the latter as more important in determining the overall value of my first-order options.

This leaves open the question whether I have to maximize along whatever value dimension I choose to stress—and, for that matter, whether I need at least to satisfice along the discounted dimension. But the present version of Rome/Riviera is meant just to flush out complexities of the real-life case. Let us replace it for a moment with a simpler case, to illustrate when satisficing could not even appear to gain a foothold. Here is a familiar sort of case that leaves no room for a reason for choice in conflict with overall value: I have to select a flight to Nice from among several options listed on the Internet. Two flights leaving at the same time match up in all other relevant respects except price: Flight A is significantly cheaper than flight B. Both are within what I can comfortably afford, but I do have other uses for the money, so A is clearly better for me than B. How can it be

rational, then, for me to satisfice, choosing B just on the grounds that it is good enough? By hypothesis, I cannot appeal to other distinguishing features of the two flights—there are none—and the case is set up to eliminate any cost in time and effort of comparing options: I have all the relevant information, and all I need to do is check a box.

In short, it does seem bizarre to go for less than the best where:

(1) the agent recognizes significant differences in value among her options,
(2) all she wants or could reasonably want is value of a single sort, or measurable along a single dimension, and
(3) she stands in roughly the same temporal relation to the different options in question.

It is (3) that allows for a different answer in bird-in-hand cases. It has to be spelled out separately because the relevance of temporal distance to rational choice need not be reflected in the value-comparison of our options, as given in (1)–(2). "Roughly" is meant to let (3) cover cases where both options are remote enough that any temporal gap between them appears insignificant.

In cases where (1)–(3) all hold, we may grant that rationality requires maximizing value. To apply this back to Rome/Riviera: on our assumptions, if I am deciding on the basis of overall value, it would be irrational for me to choose the Riviera, but if I were clear about instead prioritizing relaxation, it would be irrational for me to choose Rome. Now let us turn to a version of the case where satisficing seems perfectly sensible, even on similar assumptions about the value of my options, as given in (1)–(2).

II. Explaining Bird-in-hand Satisficing

Suppose I am already on the Riviera, enjoying a vacation that is "good enough," even by a reasonably high standard, when someone offers me a place to stay in Rome. The new offer changes the stakes. If it had come through before I made my plans, it would have made me rethink the basis for my decision. Now, however, in the middle of an enjoyable vacation on the Riviera, I can turn down the offer just on the grounds that I am happy where I am. I might grant that moving to Rome would make my vacation better—even considering the effort it would take to move, and without any doubts about the results. We also need to assume that awareness of the new opportunity does not undermine my satisfaction with my current situation: I do not now feel bored with the beach, say, when presented with the lure of a comfortable place in Rome. The point is just that, rationally speaking, a better offer is an offer I can refuse, if I already have met a certain threshold.

We may think of this case as a bird-in-hand version of Rome/Riviera. But we should note that the classification departs in some ways from the common interpretation of the saying from which the term is derived: "A bird in the hand is worth than two in the bush." This is usually taken as referring to the uncertainty of getting the two birds in the bush, but uncertainty is not what is at issue in bird-in-hand Rome/Riviera. Moreover, the saying is phrased as a claim about the greater "worth" of the bird in hand, which would mean taking temporal distance as reflected in the value-comparison of our options. By hypothesis, however, remaining on the Riviera is not represented as a better option than moving to Rome, but rather just as a rational alternative to it, even though the value-comparison runs in reverse.

Instead, something besides the value of the options we are choosing among seems to be at issue in Rome/Riviera. At a first pass, we place a certain importance on getting sufficient benefit from whatever option we choose, and a habit of continuing to assess alternatives after a choice is made—comparison shopping "after the fact," as it were—would often un-

dermine enjoyment of what we have chosen. This means that it can be rational to make ourselves relatively impervious to information about how an option fares in comparison with alternatives—even when we know that it fares less well, with any uncertainty taken into account, along with any costs in time and effort of making and executing the choice. Our assumption is that the rewards of spending time in Rome would more than make up for the bother of moving there. But it still would not be irrational for me to decline to decide on that basis, if I am happy enough where I am. I would be manifesting a generally valuable decision-making trait in my reluctance to revise settled plans.

This line of argument might just seem to set up enjoyment as a further valuable quantity to be maximized, either in place of others or added to a mix; but more fundamentally, what it appeals to is not just a pleasurable state of experience but rather a beneficial habit or pattern of choice that avoids a kind of waste. Too much after-the-fact comparison shopping would render many of our choices pointless, that is. We would be wasting the opportunity to enjoy what we have if we always were looking to trade it in for something better.

However, it would be absurd to set up waste avoidance as itself something to maximize, no matter what else is at stake. When we speak of "minimizing" waste, we just mean putting a priority on avoiding it, as part of an attempt to strike a balance between that consideration and others. The balance will indeed be something we "strike," though, in the sense of actively determining rather than simply figuring out what it is; for we have some control over where we set our threshold of satisfaction. It is partly up to us, that is, when to stop comparison shopping and thereby "content ourselves" with the status quo. But such strategies cannot be relied on to take effect—even in a case in which we do have the best option in hand. Since contentment would also improve the best option, its addition to options that are good enough will not result in a tie, so we cannot simply include the value of contentment in an optimal package of goods, or a set of equally good alternative packages, and count any option falling short of that ideal as rationally unacceptable. To be fulfillable in the normal case, the demands of rationality have to be limited to a general caution—against "too much" waste or the like—that leaves significant leeway for personal choice, on the model of an imperfect duty in ethics with a vaguely specified threshold of adequate performance, such as the duty to give a reasonable amount to charity.

The role of contentment in this argument will be explored further after we look at some alternative explanations of bird-in-hand satisficing. For the moment, let us consider some familiar cases of satisficing with regard to waste. Right after we order the item that looked most appealing on a restaurant menu, for instance, someone else's order arrives that looks even better. It is still possible to call back the waitress and switch orders; nor would a single change of mind be so embarrassing, or delay the meal all that much. But another alternative, available to most of us much of the time, is putting the comparison out of mind, to keep it from diminishing enjoyment of our own selection.

Note how the appeal to avoiding waste pulls against reliance on moderation as a virtue that entails satisficing (Slote 1989): we might also be advised to avoid waste in the restaurant case by finishing everything on our plate, even where the portions are immoderate. Intuitively, waste amounts to a kind of inefficiency, so waste avoidance exhibits an intelligible point behind satisficing that is independent of moderation and does not depend on imposing satisficing as a requirement, even in moderate but fixed degree. In another sort of case, it may make perfect sense to object on grounds of waste to throwing out electronic items every few years, when replacing them turns out to

be cheaper than repairs, but that is not to say that replacing them would be irrational. We might eventually have to conclude that it is a better option all told, but at least within a certain range of tolerable costs, we can take our choice.

An argument framed in terms of waste avoidance would not carry weight for someone committed to maximizing future benefit, on its ordinary understanding as specifiable without reference to the past. If a better opportunity comes along, such an approach would treat past choices as "sunk costs," to be dismissed from consideration. But without denying that waste avoidance can sometimes become a fetish, we can see that many of our goals have a backward-looking element. Consider, for instance, striving to make good on the promise you exhibited as a young student.

Another objection to the argument, without theoretical presuppositions, is that the point of making plans need not always be completely wasted where one trades what one has for something better. The process of carrying out a plan sometimes involves benefits independent of whether it is carried to completion. I might enjoy whatever time I spend on the Riviera, even if I do decide to cut my stay short to move to Rome. However, what is at issue in the argument is really the general habit of avoiding waste, not necessarily its effects in a given choice situation. There will be cases—replacing electronic equipment is one of them—in which giving up what one has will turn out to be advisable. But the point is just that hanging onto it also counts as rational, at any rate up to a point, in light of the *general* value of avoiding waste.

The intended view is permissive, allowing a variable role to both levels of cost-benefit assessment, general and particular. It is not meant to follow the model of two-level consequentialism, on which consequences of general rules, practices, or the like are authoritative for action. Maximizing is also permissible, remember, and of course there will also be cases where the particular assessment clearly ought to win the day. The cost of repairs may just be too large to bear.

Moreover, an exception can be made for cases promising *huge* benefits, if we think of the threshold of "good enough" as already incorporating an element of comparison, to the extent that awareness of alternatives can affect an agent's subjective satisfaction with what she has in hand. Like the poverty line, whether an option seems adequate or inadequate can vary with what else is seen as attainable. The threshold can sometimes be reset by the sort of unforeseen opportunity that makes other options "pale by comparison," as we say. But the bird-in-hand version of Rome/Riviera was set up to exclude this.

There is another version of the argument from general habits or patterns of choice that is framed in less subjective terms, though its force ultimately depends on considerations of control over attitudes, of the sort that emerged in reference to enjoyment. It appeals to practical consistency over time as a feature of rational planning agency. Part of the point of making plans is coordination—not just with others, but also among various plans of one's own (cf. esp. Bratman 1987). In bird-in-hand Rome/Riviera I may have made arrangements—to meet friends, to travel to a nearby site—that are contingent on resisting the offer to move to Rome. But the issue here is not just what works out best on a particular occasion. Achieving our aims generally depends on the sort of coherence between long-range and immediate intentions (plans and behavior) that makes us reliable planning agents.

Again, though, we need only pass a threshold: with enough consistency over time, we have the option of abandoning prior plans on occasion. In Rome/Riviera it would be well to avoid too much flitting from plan to plan, but of course there is also a point in seizing a valuable opportunity. An ideal of situational responsiveness or spontaneity also makes

sense in general terms, insofar as effective agency may be undermined by too sharp a split between immediate feelings and behavior. Consider "going through the motions" of carrying out a previous plan for which the initial enthusiasm has faded.

These two competing considerations, consistency and spontaneity, might be thought of as higher-order values, or valued traits, involving different degrees of responsiveness to the value- comparison of first-order options. They entail general patterns of choice over time, not specific choices on particular occasions, except where we have paid too little attention to one of them for too long. It is not just that several ways of balancing the opposing considerations will be tied for best. We also can sometimes defuse the conflict between them by working on our attitudes—immersing ourselves in carrying out a long-standing plan, say, to a degree that prevents comparison with alternatives and hence generates renewed interest. Once again, however, our success at such self-management tasks would seem to be too variable to support a determinate requirement as an expectation of normal agents. The most that rationality can demand, then, is striking a reasonable balance that may fall short of the ideal.

III. Some Alternatives

Put most simply, what makes bird-in-hand cases different from those cordoned off earlier in which satisficing seems irrational is the fact that in the earlier cases comparison shopping is precisely what the choice situation demands. In bird-in-hand cases it is something we can avoid, but not with the degree of control presupposed by a conception of rationality that normal agents are expected to be able to fulfill. So satisficing as here explained is

- appropriate only sometimes, and even then
- optional rather than required, but therefore
- a flexible situational response, not (or not ideally) the result of a settled character trait or virtue.

Besides the analogy to imperfect duties mentioned earlier, one might find a parallel to this account in a conception of virtue that allows for conflicts of the sort illustrated by consistency versus spontaneity. But standard virtue ethics assumes the unity of the virtues.

Now let us consider some alternatives to the present approach. We might think of these as three main proposals for replacing the kind of temporal layout implied by bird-in-hand status—essentially, the denial of (3) in the three conditions outlined earlier as making satisficing irrational: first, as we have already seen, the appeal to moderation; second, a version of the proposed explanation that appeals to multiple values as different dimensions of choice, of the sort that would count against condition (2); and third, temporal discounting, as yielding a different sort of temporal layout, with a stress on short-term costs.

First, what about cases where we have an opportunity just to augment what we have in hand, rather than replacing it with something else that we think would be better? Consider the example from Slote 1989 of turning down a snack after a satisfying lunch, even without worries about ill effects of overeating. Satisficing in such cases might also be said to involve "standing pat," but not because of a risk of undermining what we already have, as in Rome/Riviera—or for that matter, in the games from which the expression "standing pat" derives.

The case of a choice between two flights differing only in cost, in which satisficing seemed bizarre, might be seen as fitting this augmentation model, if we think of saving money on the flight as a means to having more money later. But note that an appeal to moderation might be applied in several different ways here, to yield conflicting practical directives. We have assumed that a rational agent would be moderate in what she spends

and thus choose the cheaper flight. But if exhibiting an all-round virtue of moderation is at issue, I should also be moderate in my concern with saving money. Perhaps, then, I should allow myself to choose the more expensive flight. As with reliability versus spontaneity, the answer depends on my general pattern of choice up to now—whether I have so far tended to overstress one form of moderation. But moderation, too, can be overdone. Perhaps I should also be moderate in my insistence on it.

If satisficing seems appropriate in the snack case and some others fitting the augmentation model, that might be explained by denying (2) in the list of conditions that were said to make satisficing irrational in the case of choosing a flight. Sometimes, that is, foregoing the best option serves to promote some further value besides those explicitly at issue in the choice. Moderation is one possibility, but other nonmoral virtues or valued traits may be relevant to other cases, possibly in conflict with moderation. Consider aesthetic taste: I might decline to add another object to those already on my mantelpiece in order to preserve the aesthetic balance it now exhibits, not particularly to avoid excessive display, and without being moderate in my insistence on aesthetic balance. Even in the snack case, our objection might be aimed against spoiling a balanced meal rather than excess per se.

However, this suggests a second way of questioning the proposed explanation of the rationality of satisficing. It turned on an appeal to further higher-order values—on waste avoidance, practical consistency, and so forth—that might themselves support a denial of condition (2). Consider a case of relatively immediate and cost-free replacement of an option already in hand, where nothing first-order is at issue except whether to retain that option or replace it with something better. Suppose my seatmate on the flight to Nice offers to exchange her window seat for my current seat on the aisle. I do have a preference for a window seat, though I am happy enough where I am. All I would have to do is move over, and I have no worries about accepting the favor. But I simply say, "I'm fine here." Might that be because I value something on the order of stability, or remaining settled where I am?

Our assumption is that I am not attached to any first-order feature of my current seat, even something I see as outweighed by the value of what I am offered in exchange. Though I am settled in the seat by the time the exchange is offered to me, there is nothing about it I would miss if I gave it up, in the way I might miss some distinctive features of the Riviera if I moved to Rome in bird-in-hand Rome/Riviera. There might be another version of the seat-exchange case in which, as a way of contenting myself with my assigned seat, I have dwelt at length on its advantages, most notably the fact that it makes getting out easier, to the point where it has come to have a further kind of first-order value for me, and (2) would be falsified without an appeal to higher-order values. I would then have the option of choosing on the basis of a particular aspect of overall value, compatibly with the recognition that I am turning down an option that is better overall.

Can we instead explain a choice to remain where I am as promoting some higher-order value, on the order of consistency or stability? It is we theorists, or rational critics, who typically invoke such factors, but an agent may do so herself in her capacity as self-critic, and perhaps we could also see her as motivated by them, albeit inexplicitly, in deciding what to do. However, what is important for present purposes is that this would not really amount to an alternative to bird-in-hand satisficing with regard to first-order value. For consistency or stability to be in question, the case has to exhibit a bird-in-hand temporal layout; or, in other words, we also have to deny (3) in the list of conditions making satisficing rationally unintelligible. Similarly, in Rome/

Riviera, even if my choice to stay on the Riviera is explained as an instance of a general tendency to exhibit reliable follow-through behavior, it is so only on the assumption that I either am already there or will be there at some later point when the opportunity to move to Rome comes up. In short, these cases remain cases of bird-in-hand satisficing, even if what makes them cases of *sensible* satisficing is a rationale in more general terms.

However, there is a third possibility for explaining bird-in-hand satisficing as rational in Rome/Riviera that would not apply so readily to cost-free cases like the seat-exchange case. Perhaps what is at issue is our tendency to assign lesser value to more remote goods, or temporal discounting (see Ainslie 1992). In declining to leave the Riviera for Rome, I might be seen as discounting the remote benefits of the exchange, placing more weight on its up-front costs. What deters me is the fact that moving to Rome would entail revising various plans, packing up my things, changing my expectations, and so forth, in the relatively immediate future.

Such cases might be thought of as involving a temporal trajectory with an initial hump of unpleasantness we have to get over in order to get where we otherwise would like to go. Even where I am convinced that the eventual benefits of change will *outweigh* its near-term costs, I might balk at having to go through a period of concentrated costs in order to accrue those benefits.

Temporal discounting is something we find rationally intelligible, even if not ideal, and it might seem to be justified by a version of a common objection to simple maximizing views: that the order of events matters. For instance, temporal order matters—though in the opposite direction—when we assess whether a life is happy: we place more importance on how things work out in the end. The satisfaction of achieving one's central goals in life is thought to justify an earlier period of stress and deprivation—in contrast both to staying at an average level and to unrestrained enjoyment in one's youth followed by later regret and self-reproach, while maintaining the same overall quantities of value and disvalue. Perhaps what is at issue in bird-in-hand cases, then, is a heavier weight on relatively immediate disvalue when it comes to smaller-scale choices.

In the cases featured here, however, distance in time apparently does not affect the agent's estimate of the *value* of her options, but just her preference for one over the other. In bird-in-hand Rome/Riviera, I do acknowledge that moving to Rome would be better overall, and in the seat-exchange case—if we can allow for it by stressing the very slight effort it would take to move over—I still have a preference for a window seat, despite disclaiming any need for one. But a plausible application of temporal discounting would essentially impose a limit on the short-term costs we ought to bear.

Temporal discounting would not replace bird-in-hand satisficing, though, if our present level of satisfaction serves to set the threshold of intolerable costs. Declining to correct a *bad* situation because of qualms about what that would require in immediate terms would be weak-willed, except where the costs are high. It is when things would otherwise be good enough that it seems sensible to avoid a short-term ordeal. When the present time is the benchmark, we are advised to "leave well enough alone."

However, the explanation in terms of tolerable costs might instead be reframed in a way that does not privilege the present time but again appeals to a further dimension of value: the value we place on maintaining a reasonably stable level of comfort, say. This would count against the upheaval of moving in Rome/Riviera—but also, more generally, against undergoing any period of concentrated costs, whether up-front or ultimate (as in reversals of fortune and downhill slides). It would also explain why bird-in-hand satis-

ficing seems no less rational in prospect—in cases of serial decision making—than it does from the present standpoint.

When we do put up with an ordeal, though, and things work out for the better, we often look back on it as justified. So bird-in-hand satisficing seems to allow for a shift with temporal standpoint; but it is not quite the same as the shift predicted by temporal discounting, which also would take *bad* times later to be justified by sufficient enjoyment in the near term. However, it seems rationally permissible (at a minimum) to decline an immediately appealing option that entails a large cluster of remote-term costs, even if we impose a higher threshold in estimating when such costs are intolerable. In short, then, insofar as temporal discounting entails overestimation of short-term *value*, and not just disvalue, its treatment of cases is intuitively problematic in ways not entailed by bird-in-hand satisficing.

IV. The Present as First-Person Benchmark

The overlap with temporal discounting in the treatment of short-term disvalue may still bring out a feature that sets bird-in-hand cases apart from cases of satisficing as conceived by decision theorists. Placing more weight on upfront costs than on costs at some later time may be rational for an agent, but not for someone in the position of an external observer advising her on what is to her advantage. By the same token, an objective assessment of an option as meeting a threshold of adequacy presumably could be made prior to having it in hand and would be unaffected by later changes in the agent's deliberative standpoint, whereas the rationality of bird-in-hand satisficing depends on an assessment of options from the standpoint of the agent at a certain point in time.

What counts as good enough from the agent's standpoint may change over time, as indicated earlier by the observation that an offer may sometimes reset the threshold by lowering the agent's assessment of her other options. Similarly, once a better option is in hand, an earlier state of affairs may no longer be good enough. One might acknowledge that it was or would have been good enough, while ceasing to see it as good enough now. Once in Rome, perhaps I could no longer be satisfied with the Riviera. The threshold of adequacy can shift over time, then, in ways that bear on rational choice but need not be reflected in the judgments of a third party, even an adviser with the agent's interests at heart, but assessing them from an external standpoint.

This may be part of the reason why bird-in-hand satisficing does not fall subject to an argument in recent consequentialist literature (see Bradley 2006) that cuts against a blanket permission to satisfice—in the first instance, with regard to the total good, though the argument also is meant to apply to rational satisficing. It turns on the assumption that the threshold of "good enough" might be lower than the level already reached. In the moral case, it would be wrong to cause gratuitous harm to others, even if it still leaves them sufficiently well-off. By the same token, gratuitous self-sacrifice would seem to be irrational—giving up something good that is already in hand just because one could have done well enough without it. It is another question whether one still could do well enough without it.

In any case, the argument would not rule out leaving things as they are, assuming one is doing well enough now. This does not apply, though, to situations with a tendency to deteriorate without further action on the part of the agent. A situation that will become bad if I do nothing (my level of physical fitness, say) does not really count as a bird in hand, even if I am satisfied with it for the moment. Where I do have a bird *firmly* in hand, the form of satisficing defended in this paper will allow me to turn down the offer of something

better—but not to let go of what I have without a reason.

The significance of the agent's temporal standpoint might seem to make the argument for satisficing depend on imperfections in human foresight. Besides not always being able to tell whether what I have is really all I need, I cannot always predict that my contentment with it will continue—given the limitations on our ability to control such attitudes that were featured in the proposed explanation of bird-in-hand satisficing. A requirement to strike the best possible balance among the multiple values at issue in a given choice was rejected as unrealistic, particularly in light of their inclusion of subjective states like contentment that are subject to only limited control. However, this is still an "in principle" justification of satisficing, despite its restriction to agents with human capacities—for whom rationality in the intended sense is supposed to be the normal condition.

A requirement to maximize would make excessive demands, that is, given the role of contentment as simultaneously a good state of feeling and a valuable behavioral propensity, to be unreceptive to comparison. An agent engaged in serial choice can never conclude with certainty that nothing better will turn up later, but remaining open to further possibilities would always have a tendency to undermine the value of what she has in hand, even where it does turn out to be the best option. We need to be able to get ourselves at least to stop comparison shopping at this level, and for agents like us the way to do so is by cultivating a habit likely to be exercised before we get there.

Various "stopping rules" might be formulated in support of maximizing—advising us, for instance, to stop scanning for better options only once we have determined by some reasonable procedure that we have the best option in hand, or only after a certain amount of time has passed with no better option appearing. But these would seem to reinforce the tendency to keep comparing options—while demanding that comparison somehow be set to stop abruptly when a certain comparative criterion is met. After that, but only then, we must make ourselves impervious even to further comparative information that arises without effort on our part. There is an obvious tension here, however. Trying to force adherence to a rule—a rule banning further comparison, at a point reached precisely *by* comparison—would more likely lead to frustration than contentment in a creature that operates on the basis of general habits.

Instead, the way to stop comparing options and end up feeling content is to cultivate a focus on aspects of the current situation that are satisfying in themselves. But a habit of focusing contentedly on what one has will naturally extend to nonideal cases. We would be well advised to stop scanning for better options once we have an option in hand that counts as good enough on some noncomparative criterion. This is not to say that the threshold of adequacy cannot in fact be set or reset by comparison. It may be a function of the option we take as best, either from the outset (cf. Hurka 1990) or in those cases where we cannot help being dissatisfied with what we have in light of a particularly tempting offer. But to the extent possible, once we have a satisfying option in hand, it would be well to avoid a comparative frame of mind.

Reasonable advice would be to make our demands early on, when temporally equidistant or remote options are presented to us for comparison, but then adopt a different outlook when we have something in hand: "Aim at the best, but be content with good enough." But such advice is at odds with standard decision theory if it endorses our common tendency to treat the loss of something as worse than a failure to gain some equivalent benefit. Common errors in reasoning of the sort emphasized in work on heuristics as departures from decision-theoretical rationality turn on the fact that describing something as a

loss rather than a benefit—citing the number of deaths rather than lives saved if a drug is administered, say (see Tversky and Kahneman 1981)—makes people more averse to it. However, we are now in a position to see that there is more to bird-in-hand satisficing than a simple attachment to what we have. Instead, it involves the exercise of a general habit that is needed to make the most of what we have.

The conclusion of this paper is not just that exercising such habits leads us to approximate optimal decisions overall, in the manner of a heuristic. Attitudes on the order of contentment count as part of what it is rational for us to choose in particular cases, rather than just as general guides or inducements to choice assessed as rational in other terms. The general point of the paper is that we choose subjective attitudes as well as acts, with important effects on the value of our options, but with variable control that justifies protecting ourselves by inducing a habit of contentment with options that may fall short of the ideal.

The defense here of bird-in-hand satisficing is not meant to displace decision theory—there still is a point in a mathematizable system that treats rationality as an objective property of acts—but just to put the burden on those who would apply it to everyday choice to show how a requirement to maximize could be put into effect without undermining itself by generating discontent. The argument appealed to the way we need to balance conflicting values pertaining to habits or patterns of choice—ultimately, those involved in comparing options versus contentment with satisfying aspects of the option we have in hand. It was based on human psychological limitations, though not limitations on information processing—and not *imperfections*—unless one counts the fact that our knowledge that an option is best is not generated automatically, in the manner of a machine, but depends on the exercise of habits of mind involved in seeking information. Time came into the argument, not because it is too short for optimal decision making, but because the assessment of our options for decision varies quite properly with their temporal layout and our own standpoint in time.

University of Maryland, College Park

NOTE

Thanks are due to Karen Jones, Samuel Kerstein, Christopher Morris, and participants in a Philosophy Department colloquium at Florida State University for comments on earlier drafts of this paper. I also benefited from later discussion of some issues relevant to my final section by Douglas Portmore, Ben Bradley, and other contributors to http://peasoup.typepad.com/peasoup/.

REFERENCES

Ainslie, George. 1992. *Picoeconomics: The Strategic Interaction of Successive Motivational States within the Person* (Cambridge: Cambridge University Press).

Bradley, Ben. 2006. "Against Satisficing Consequentialism," *Utilitas*, vol. 18, pp. 97–108. Bratman, Michael E. 1987. *Intention, Plans, and Practical Reason* (Cambridge, Mass.: Harvard University Press).

Gigerenzer, Gerd, and Reinhard Selten, eds. 2002. *Bounded Rationality: The Adaptive Toolbox* (Cambridge, Mass.: MIT Press).

Greenspan, Patricia S. 2005. "Asymmetrical Practical Reasons," in *Experience and Analysis: Proceedings of the 27th International Wittgenstein Symposium*, ed. M. E. Reicher and J. C. Marek (Vienna: oebv&hpt), pp. 115–122.

_____. 2007. "Practical Reasons and Moral 'Ought,'" *Oxford Studies in Metaethics*, vol. 2 (Oxford: Oxford University Press), pp. 172–194.

Hurka, Thomas. 1990. "Two Kinds of Satisficing," *Philosophical Studies*, vol. 59, pp. 107–111.

Simon, Herbert A. 1955. "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics*, vol. 69, pp. 99–118.

Slote, Michael. 1985. *Common-Sense Morality and Consequentialism* (London: Routledge & Kegan Paul).

_____. 1989. *Beyond Optimizing: A Study of Rational Choice* (Cambridge, Mass.: Harvard University Press).

Tversky, Amos, and Daniel Kahneman. 1981. "The Framing of Decisions and the Psychology of Choice," *Science*, n.s., vol. 211, pp. 453–458.