

Satisficing and Maximizing
Moral Theorists on Practical Reason

Edited by
MICHAEL BYRON
Kent State University



PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa
<http://www.cambridge.org>

© Cambridge University Press 2004

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2004

Printed in the United States of America

Typeface ITC New Baskerville 10/13 pt. *System* L^AT_EX 2_ε [TB]

A catalog record for this book is available from the British Library.

Library of Congress Cataloging in Publication Data

Satisficing and maximizing : moral theorists on practical reason / edited by Michael Byron.
p. cm.

Includes bibliographical references and index.

ISBN 0-521-81149-X – ISBN 0-521-01005-5 (pbk.)

1. Decision making – Moral and ethical aspects. 2. Practical reason. I. Byron,
Michael, 1964–

BJ1419.S28 2004
170'.42–dc22 2004043584

ISBN 0 521 81149 X hardback
ISBN 0 521 01005 5 paperback

Contents

<i>Contributors</i>	<i>page</i> vii
Introduction	1
<i>Michael Byron</i>	
1 Two Views of Satisficing	14
<i>Michael Slote</i>	
2 Satisficing as a Humanly Rational Strategy	30
<i>David Schmidtz</i>	
3 Maxificing: Life on a Budget; or, If You Would Maximize, Then Satisfice!	59
<i>Jan Narveson</i>	
4 Satisficing and Substantive Values	71
<i>Thomas Hurka</i>	
5 A New Defense of Satisficing	77
<i>Michael Weber</i>	
6 Satisficing: Not Good Enough	106
<i>Henry S. Richardson</i>	
7 Why Ethical Satisficing Makes Sense and Rational Satisficing Doesn't	131
<i>James Dreier</i>	
8 The Plausibility of Satisficing and the Role of Good in Ordinary Thought	155
<i>Mark van Roojen</i>	
9 Satisficing and Perfectionism in Virtue Ethics	176
<i>Christine Swanton</i>	

10	Could Aristotle Satisfice?	190
	<i>Michael Byron</i>	
11	How Do Economists Think About Rationality?	213
	<i>Tyler Cowen</i>	
	<i>Bibliography</i>	237
	<i>Index</i>	243

Contributors

Michael Byron is Associate Professor of Philosophy at Kent State University, where he has been teaching since 1997. His research interests include ethical theory, rational choice theory, and the history of ethics; he is the co-author (with Deborah Barnbaum) of *Research Ethics: Text and Readings*, published by Prentice-Hall, and articles in ethical theory and theory of rationality.

Tyler Cowen is Holbert C. Harris Professor of Economics at George Mason University. He has published extensively in economics and philosophy journals, including *American Economic Review*, *Journal of Political Economy*, *Ethics*, and *Philosophy and Public Affairs*. His *In Praise of Commercial Culture* and *What Price Fame?* were published by Harvard University Press, and his latest book, *Creative Destruction: How Globalization Is Changing the World's Cultures*, was published in 2003 by Princeton University Press. He is currently writing a book on the concept of civilization and its relation to political philosophy.

James Dreier is Professor of Philosophy at Brown University. He has been a visiting lecturer at Monash University and John Harsanyi Fellow at the Social and Political Theory program at the Research School of Social Science at the Australian National University. His main work is in meta-ethics and practical reason.

Thomas Hurka is Jackman Distinguished Chair in Philosophical Studies at the University of Toronto. The author of *Perfectionism* (1993), *Principles: Short Essays on Ethics* (1993), and *Virtue, Vice, and Value* (2001), he works

primarily on perfectionist views in ethics, political philosophy, and the history of ethics.

Jan Narveson is Professor of Philosophy at the University of Waterloo in Ontario, Canada. He is the author of more than two hundred papers in philosophical periodicals and anthologies, mainly on ethical theory and practice, and of five published books: *Morality and Utility* (1967), *The Libertarian Idea* (1989), *Moral Matters*, (1993; 2nd ed. 1999) and *Respecting Persons in Theory and Practice* (2002); and, with Marilyn Friedman, *Political Correctness* (1995).

Henry S. Richardson is Professor of Philosophy at Georgetown University. He is the author of *Practical Reasoning about Final Ends* (Cambridge University Press, 1994) and *Democratic Autonomy* (Oxford University Press, 2002) and co-editor of *Liberalism and the Good* (Routledge, 1990) and *The Philosophy of Rawls* (5 vols., Garland, 1999). Building on "Beyond Good and Right: Toward a Constructive Ethical Pragmatism," *Philosophy & Public Affairs* 24 (1995), he is working to articulate a moral theory centering on non-maximizing, non-satisficing modes of moral reasoning.

David Schmidtz is Professor of Philosophy and joint Professor of Economics at the University of Arizona. His *Rational Choice and Moral Agency* (Princeton University Press, 1995) expands upon his chapter reprinted in this volume. He co-edited *Environmental Ethics: What Really Matters, What Really Works* (Oxford University Press, 2002) with Elizabeth Willott and co-authored *Social Welfare and Individual Responsibility* (Cambridge University Press, 1998) with Robert Goodin. His current projects are *The Elements of Justice* (Cambridge University Press, 2005) and *The Purpose of Moral Theory*.

Michael Slote is UST Professor of Ethics and Professor of Philosophy at the University of Miami, Coral Gables. He was previously Professor and Chair of the philosophy department at the University of Maryland, and before that he was Professor of Philosophy and a Fellow of Trinity College, Dublin. The author of many articles and several books on ethical theory, he is currently completing a large-scale study of "moral sentimentalism." He is a member of the Royal Irish Academy and a former Tanner lecturer.

Christine Swanton teaches at the University of Auckland, New Zealand. Her field is virtue ethics, in which she has recently published a book with Oxford University Press, entitled *Virtue Ethics: A Pluralistic View*. She is

currently working in role ethics, Nietzschean virtue ethics, and Humean virtue ethics.

Mark van Roojen is Associate Professor of Philosophy at the University of Nebraska–Lincoln. He has taught philosophy as a visitor at Brown University and at the University of Arizona. His main research interests are in meta-ethics, ethics, and political philosophy.

Michael Weber is Assistant Professor of Philosophy at Yale University. His research and teaching are in moral and political philosophy, focusing on practical reason, rational choice theory, and ethics and the emotions. His papers have appeared in *Ethics*, *Philosophical Studies*, and the *Canadian Journal of Philosophy*.

Introduction

Michael Byron

It is testimony to the breadth of thought of Herbert Simon, the man who conceived the idea of ‘satisficing’, that the concept has influenced such a wide variety of disciplines. To name a few: Computer science, game theory, economics, political science, evolutionary biology, and philosophy have all been enriched by reflection on the contrast between choosing what is satisfactory and choosing what is best. Indeed, these disciplines have cross-fertilized one another through the concept. So one finds satisficing computer models of evolutionary development, satisficing economic models of international relations, satisficing applications of game theory within economics, and philosophical accounts of all of these.

Philosophical interest in the concept of satisficing itself represents a convergence. The fecund and appealing idea of choosing what is satisfactory finds a place in the theory of practical reason, or thinking about what to do. The appeal of the concept derives partly from the fact that what is satisfactory is, well, satisfying. Satisfaction is generally good, and goods of this generality feature prominently in any account of practical reason. More noteworthy is the fact that the concept of satisficing finds application from so many perspectives, even within the relatively narrow confines of moral theory.

In any conversation of this complexity it is always in point to ask whether the participants are talking about the same thing. So of course this issue arises with respect to the essays collected in this volume. I will not try to resolve that issue here; instead, I would like to explore the starting points that have led the authors to their views about satisficing. As so often happens in any theoretical enterprise, where one begins has

a crucial – and often decisive – impact on where one ends up, at least if one’s conclusion follows from one’s premises.

Simon Says

Among the first statements of his account of satisficing is Simon’s essay “A Behavioral Model of Rational Choice.”¹ Several contributions in this volume summarize Simon’s argument, so I will be brief. After characterizing the choice situation and defining his terms, Simon contends that in typical choice situations the application of classic rules of choice such as “max-min,” the “probabilistic rule,” and the “certainty rule” involves calculative and cognitive skills that no human being possesses. Maximizing expected utility – or choosing one’s actions so that they are most likely to bring about states of affairs that one prefers – seems out of reach for people like us. Simplification seems to be in order.

Simon simplifies rational choice along two dimensions. One is the value function. Maximization requires rational agents to assign a utility, or numerical index of preference, to each possible outcome of every available alternative action. Everything that might happen after one acts must be rated on a common numerical scale. Putting the point this way already seems daunting. Simon’s model allows agents instead to use a simpler evaluation function in two values, (1,0), corresponding to “satisfactory” and “unsatisfactory.”² So rather than, for example, having to decide exactly how much better two heads are than one, we might simply say that both are satisfactory, and zero heads would be unsatisfactory.

The second dimension of simplification appears in Simon’s satisficing rule itself. According to that rule, rationality requires an agent first to identify the set of all satisfactory outcomes of the choice situation, and then to choose an alternative all of whose outcomes are in the set of satisfactory outcomes. More briefly, it’s rational to choose any action that guarantees a satisfactory outcome. That way, whatever happens, one will be satisfied. This rule is simpler than maximizing in virtue of eliminating probabilities from rational choice. For when I maximize, I must weight the utility of each possible outcome by the probability that it will occur.

An example might help make the simplification more evident, mainly by exposing the complexity of maximizing. Suppose that I am at the racetrack planning to make a bet. I am prepared to make one of three choices: no bet, a \$10 bet on Toodle-oo, or a \$10 bet on Beetlebaum. The guaranteed outcome of no bet is that I lose nothing and keep my

\$10. The odds on Toodle-oo are 4:1, which means the payoff if the horse wins would be \$50 (\$40 plus the \$10 bet). The odds also indicate that in the oddsmakers' judgment the horse has a 0.2 probability or 20 percent chance of winning. Beetlebaum, on the other hand, is the long shot, paying 24:1. My \$10 bet would yield \$250 if Beetlebaum won, but there's only a 0.04 probability or 4 percent chance of that. Let's call these options *N* (no bet), *T* (\$10 on Toodle-oo), and *B* (\$10 on Beetlebaum).

We can then calculate the expected utility of each of these options according to the following formula:

$$\begin{aligned} EU(A) = & [P(O_1/A) \times U(O_1)] + [P(O_2/A) \times U(O_2)] \\ & + \dots + [P(O_i/A) \times U(O_i)] \end{aligned}$$

Each of the O_i represents a possible outcome of the act *A*, $U(O_i)$ is the utility of that outcome, and $P(O_i/A)$ is the probability of O_i given *A*.³ We can thus calculate the expected utilities of each of our three acts. Let's assume that utility is convertible with dollars.

$$\begin{aligned} EU(N) &= P(\text{no loss}) \times U(\text{no loss}) = 1 \times 10 = 10 \\ EU(T) &= [P(T \text{ wins}) \times U(T \text{ wins})] + [P(T \text{ loses}) \times U(T \text{ loses})] \\ &= [0.2 \times 50] + [0.8 \times -10] \\ &= 10 - 8 = 2 \\ EU(B) &= [P(B \text{ wins}) \times U(B \text{ wins})] + [P(B \text{ loses}) \times U(B \text{ loses})] \\ &= [0.04 \times 250] + [0.96 \times -10] \\ &= 10 - 9.6 = 0.4 \end{aligned}$$

Clearly, if that's all there is to the choice situation, then the maximizing choice is not to bet. But I have probably omitted an element of the second two alternatives, namely the enjoyment of betting. Suppose that the excitement alone is worth \$10 to me; in that case, I should add 10 to the expected utilities of *T* and *B* (but not to *N*, where I make no bet), and the bet on Toodle-oo emerges as the best choice. In any case, the complexity of the calculation – even for this simplistic example – is evident.

Now contrast Simon's satisficing approach. First, I rate the outcomes as satisfactory or unsatisfactory. Suppose I rate as satisfactory winning more than \$100 or losing \$10 (because that would mean I had had a chance to win); no loss or gain (the outcome of not betting) is an unsatisfactory outcome. In that case, *N* would be irrational according to the satisficing rule, because it would guarantee an unsatisfactory outcome. *T* would also be eliminated by the rule, because it has a possible outcome that is unsatisfactory (namely, winning \$50). Hence, the satisficing rule identifies *B*

as the rational choice, because all of its possible outcomes are satisfactory. Notice that I have assigned no utilities or probabilities and performed no calculations. I have, of course, evaluated the outcomes, but only to the extent that I have rated them as “satisfactory” or “unsatisfactory.”

Simon’s satisficing rule can also function in another way. Suppose you find yourself searching for alternatives because they are not all in view. In that case, it is impossible to apply expected utility theory to all the alternatives, because they aren’t known. Simon’s satisficing rule can be employed as a “stopping rule” in this kind of choice situation. That is, it can provide a principled way to stop searching for alternatives. So if you were looking for a suitable wine to serve with dinner, you might adopt as a rule the idea of stopping your search upon finding a satisfactory wine – one that is “good enough.” This approach is consistent with the other sort of satisficing, inasmuch as you use a simplified valuation function (satisfactory, unsatisfactory) instead of assigning utilities to every possibility, and you choose a course of action that is guaranteed to be satisfactory.

The computational tractability and theoretical simplicity of satisficing are its most attractive features for Simon. And, at least initially, he presents the model of rationality as both normatively and descriptively more adequate for human beings than a maximizing conception. It is more normatively adequate in making rationality feasible for creatures like us. How could we be required to maximize, when in most cases we cannot? A maximizing conception of practical rationality entails that virtually all of our actions are irrational, and that consequence seems counterintuitive. Yet defenders of maximizing insist that their conception sets a standard against which human choices can be judged: Our choices are rational *to the extent that* they approximate the ideal established by the maximizing conception. Moreover, maximizing theories are most defensible when they incorporate constraints (time, money, cognitive resources, etc.). Such theories propose as the normative standard not bare maximizing, but maximizing within given constraints. This debate has implications for maximizing conceptions of rationality, such as rational choice theory, because if a satisficing model is normatively superior, then rational choice theory loses its claim to be the best account of practical reason.

On the other hand, satisficing models can claim to be descriptively adequate, as when they contend to be better accounts of the way people actually choose. Does anyone actually rate all the possible outcomes along a single scale of utility? Don’t many of us think, upon reaching a particular outcome, “Well, that’ll do”? As many of the contributions will remind us,

satisficing is rational as a time- and other resource-saving strategy: Given our limited resources, we sometimes settle for what's good enough in order to devote resources elsewhere. We could hold out for the best price when buying or selling a car, but that could consume a lot of time and energy that we would prefer to spend elsewhere. And so we take an offer that is good enough. Defenders of maximizing have a response to this claim of descriptive superiority; I'll return to it later.

Supererogation

The concept of supererogation has for many years been at home in "common sense morality." Meaning "above what is required," supererogatory acts exceed some threshold, typically a threshold of moral duty. Heroic and saintly acts are often regarded as "above and beyond the call of duty," and they are paradigm examples of morally supererogatory types of action. Morality might require me to assist a stranger in an emergency if the situation presents little or no cost or risk to me, for example by calling 911. But I am probably not morally required to risk my life to assist someone I don't know. Not required, but if I choose to help anyway, that's something especially admirable, something supererogatory.

Michael Slote has argued that the concept of supererogation can be applied analogously in the context of practical reasoning.⁴ Just as we are often not required to do the *morally* best action, so we are often not required to do the *rationality* best action. Slote goes to considerable lengths to spell out the analogy, arguing by reference to a range of cases which he finds intuitively satisfying that in many contexts one might rationally decline to do an action that one judged better than the action one in fact chose. His satisficing conception of rationality emerges as a competitor to a maximizing account in virtue of capturing the appealing idea that rationality does not always require us to do and choose the very best of everything.

An interesting feature that emerges from Slote's view – and a point of contention among defenders of the satisficing conception of rationality – is that it can be rational to choose an option that one judges to be *inferior*. Suppose I prefer Zinfandel to Shiraz, and suppose I have a bottle of each. Other things equal, one might expect that I would choose the Zinfandel, in virtue of the fact that I prefer it. Yet Slote contends that I need not – that rationality does not require me to, that it would not be irrational of me not to – choose the bottle that I prefer. If the Shiraz is satisfactory, then I can rationally choose it.

To see why this position is puzzling, remember that when I say, “I prefer Zinfandel,” that means I prefer it at the moment of choice, and so I’m not indifferent between the two. I might prefer Shiraz under some circumstances, but those don’t apply (or I wouldn’t have the preference that I do have at the moment of choice). The concept of preference is closely linked to choice: Ordinarily, many theorists suppose, preference determines rational choice, so that to prefer something is to be disposed to choose it when the time comes. True, the Shiraz is satisfactory; but, according to the story, I *prefer* the Zinfandel. My preference presumably captures all the reasons that I have for choice. Tote them all up, and I like the Zinfandel better. So how could it be rational to choose the Shiraz? The fact that it is satisfactory doesn’t seem to do the trick.

But things are not necessarily as bad as they might seem for Slote and others who defend this kind of strong view of satisficing. The line of thinking that problematizes satisficing depends on a particular conception of ‘preference’ and an assessment of actions in terms of their outcomes. This thinking is, crucially, consequentialist, because it evaluates actions by how well they succeed in bringing about preferred consequences. Yet, as the discussion of supererogation might suggest, it is possible to think of rationality more in terms of duties than consequences. The concept of supererogation – of actions above and beyond what duty requires – is at home in deontology, not consequentialism, which is typically understood as a family of theories embodying a maximizing conception of the good. If right action is the best, how could one go “above and beyond” that? We might understand rationality to impose certain cognitive and practical duties on us, duties that are related to our preferences but not necessarily cashed out in consequentialist terms. Our rational duties might thus impose a threshold on our actions, such that if we fail to meet or exceed that threshold we act irrationally. Yet it does not follow, on this view, that rationality demands always seeking the most preferred outcome. Our duties need not be that stringent – indeed, many theorists argue that they are not so for moral duties. The details of this sort of view would have to be spelled out; some of the contributions in this volume approach, in interestingly different ways, this task.

Moderation

An idea that appears several times in this volume is the notion that the adoption of satisficing as a strategy of rational choice exhibits a virtue, especially the virtue of moderation. This point seems to apply a kind

of moral evaluation to rational choice. Maximizing and optimizing can seem greedy: Those who maximize are by definition always seeking more, indeed as much as possible. Misers maximize their money, gluttons maximize their food, sadomasochists maximize pain, and hedonists maximize pleasure. In each of these cases, and perhaps generally, maximizing appears to be morally objectionable. We need not think it is maximizing alone that is morally objectionable in these vices; gluttony and greed might be wrong for other reasons as well. But they each involve maximizing, and some theorists claim that this feature contributes to their wrongness.

In contrast, moderation has from antiquity been regarded a virtue. True, theorists have defined the term differently – Aristotle’s *sophrosune* is distinct from the moderation of genteel British society praised by Victorian novelists, for example. Yet when contrasted with maximizing, the points of overlap among these concepts might be more significant than their differences. For essential to any concept of moderation is the idea of steering between excess and deficiency: neither too much nor too little. And to have application, the idea of avoiding excess must generally eschew maximizing. Those who pursue moderation might thus be led to embrace a model of practical rationality given in other than maximizing terms.

Moderate folks – or anyone else who pursues the virtue of moderation – might prefer to understand practical rationality in terms of satisficing rather than maximizing. For in satisficing one need not always seek the best or the most. To seek what is good enough – especially when the best is an option – might emerge as an expression of the virtue of moderation. In the cafeteria line, I might see that I could take three pieces of chocolate cake and maximize my enjoyment. But I might on reflection decide that three pieces are too many, and that one is enough and so the rational choice. This kind of deliberation embodies the intuitive appeal of a satisficing account of rationality for proponents of the virtue of moderation.

Notice that the form of argument here is distinct from that linking supererogation with satisficing. There, it was an analogy that provided the conceptual link. Rationality is supposed to be akin to morality in imposing practical duties. These duties are such that some actions meet them, others exceed them, and still others fall short of meeting them. Maximizing is analogous to morally supererogatory actions in exceeding the duties imposed by rationality. Satisficing is akin to morally permissible actions that are not heroic or otherwise supererogatory: In satisficing we

fulfill our rational obligations while recognizing that it is possible to be “super-rational” and maximize. The argument here, in contrast, is that the pursuit or exercise of a virtue rationally leads one to choose in a distinctive way. The moderate choice is a satisficing choice; and in general a satisficing conception of practical rationality seems more moderate. Satisficing emerges as, if not itself a virtue, a strategy of choice that we might often use in the service of or to express a virtue.

Not everyone finds this line of thinking persuasive. David Schmidtz, for one, has challenged the notion that moderation and satisficing exhibit any interesting conceptual connection (see Chapter 2). He points out that the contrast class of the moderate is the immoderate, not the maximizing. The satisficer as such is satisfied with a certain bundle of goods, but nothing in the idea of satisficing ensures that the satisfactory bundle is moderate. I might find three dozen cookies a satisfactory serving, but my doing so does not make that quantity moderate. Similarly, moderation need not be satisfactory in every case. Moreover, the maximizer could end up with a moderate amount, if time and other resource limitations make further pursuit of the good at stake too costly.

Finally, we might note that the form of argument here is odd. Why should theorizing about practical rationality be responsive to moral virtues like moderation? Some theorists would certainly approve of this connection: Those who discover a substantial connection between rationality and morality – like a Kantian who identifies practical rationality as the criterion of morality – would insist that any adequate account of practical reasoning will have to end up endorsing all (and perhaps only) moral actions. Such an approach might be correct, but it seems to beg the question against instrumentalists, who often defend maximizing conceptions of practical rationality. That is, the defender of satisficing claims that employing the strategy is rational in virtue of its expressing certain virtues, whereas an instrumentalist might challenge the rationality of “virtuous” action that fails to maximize (“If you’re so smart, why ain’t you rich?”). This approach to defending satisficing thus seems at home in a larger defense of a more substantive conception of practical reasoning, one that links rationality and morality in substantive ways.

Consequentialism

In an early paper on satisficing, Michael Slote challenged the intuitive connection between consequentialism and maximizing.⁵ From James Mill and Jeremy Bentham, utilitarians and other consequentialists have

embraced a maximizing conception of right action. John Stuart Mill's principle of utility, for example, declares an action right "in proportion as it tends to promote happiness," or pleasure and the absence of pain.⁶ It is easy to see the appeal of such a view, once we recognize that the fundamental insight of consequentialism is to make the world a better place. Given a choice, it makes sense to choose the best. Once we make the consequentialist move and declare that the only features of actions relevant to moral evaluation are their consequences, we seem to have every reason to strive to bring about the best consequences we can, whether 'best' is cashed out in terms of pleasure, preference satisfaction, or agent-neutral value.

Slote presents satisficing consequentialism as an alternative conception of morality. His argument claims several virtues for the conception, including the capacity to account for supererogation. It's difficult to see how a maximizing conception of morality can allow room for supererogation: If I'm required in every case to choose the best available option, how could I ever do *more* than morality required? What would supererogation mean in a maximizing context? So a possible objection to maximizing consequentialism is that it leaves no conceptual room in moral theory for supererogation. Satisficing consequentialism allows for supererogation by, in principle, leaving a gap between a morally permissible action that is good enough and one that is the best, allowing that in some cases the two might coincide.

Notice that the earlier section on supererogation focused on the idea of rational, rather than moral, supererogation. There, we came to the idea of satisficing through the theory of rationality, and the point was to introduce the concept of rational supererogation and to build a theory around that idea. Here, the concept of moral supererogation is supposed to be intuitive, and we entertain the idea of an alternative to traditional maximizing consequentialism in order to account for moral supererogation.

It is, of course, open to defenders of a more traditional consequentialism to jettison the concept of supererogation. Though intuitive, the idea of supererogation does not hook up well with the idea of rational choice. Given a choice between A and B, if A is better, why choose B? The point is especially pressing if the values are moral: If A is morally better, on what moral ground could one choose B? Defenders of maximizing consequentialism might point out that the notion of supererogation is most at home in a deontological theory, where the concept of duty plays a significant role. Once the idea of duty is in place, it's easier to make

sense of how an action can be “above and beyond” duty. Because consequentialism has traditionally found less place for the idea of duty (and deontology’s correlative evaluation of action in terms of motivation), it might make sense to resist the idea that supererogation ought to play a significant role in consequentialist thought. If that’s right, part of the intuitive motivation for satisficing consequentialism goes away.

Incommensurability

Here’s a general approach to arguing that satisficing is not a distinctive choice strategy, but rather just one kind of optimizing strategy. First ask: In virtue of what is an alternative “good enough”? The satisficer as such chooses an alternative because it is, in some way, good enough, whether or not it is the best. Assume that doing so is rational, in some sense. But something about the alternative must rationalize or justify the choice: It is presumably some feature of the alternative that makes it good enough. However the chooser answers this question, the feature(s) mentioned can be built into a conception of good, utility, or whatever according to which the choice is optimizing. Or so holds this line of thought.

This strategy takes advantage of the conceptual connection discussed earlier between preference and choice. Where preference is clear, it seems to determine choice. What would it mean to prefer A over B but to choose B? Such choice is surely *possible*: One might be under a spell, or in the grip of a passion, or otherwise impaired and prevented from executing a rational choice. But such instances are hardly paradigms of rationality. For the choice of B over A to be rational, it must be superior to A in some respect, or at least equal to A on balance. One’s initial description of the choice situation might be inadequate or incomplete in some respect relevant to understanding the choice as rational. So, the proponent of this view will argue, once we take into account the entire picture, B emerges as the maximizing choice.⁷

For example, suppose that I decide to buy a certain model of car to drive to work, and suppose I do so because it is “good enough.” What makes it good enough? I might value the reliability of the car, its purchase price being within a certain range, the style and comfort, and so on. In all of these respects, the model I settle on is satisfactory. Now, if we understand my preferences that pertain to the car purchase to include budgetary preferences (both time and money), then it’s possible to portray this choice as maximizing over all of my preferences. Sure, the car is “good enough” with respect to reliability and style. But I do not wish to

spend any more time or money making this purchase than I must, and so within those constraints the choice of the first model with satisfactory features along other dimensions emerges as the best overall. That's how this view collapses all satisficing into optimizing: If rational choice is to be intimately linked to an account of rational preference satisfaction, then satisficing is rational only if optimific. The present instance illustrates the general strategy for reducing satisficing to a kind of optimizing.⁸

One obvious way to resist this kind of strategy would be to claim that the rational preference account is a mistake. That kind of account presupposes that all values are commensurable, so that we might always weigh any two alternatives against each other pairwise to determine which is better. In the face of value incommensurability, it will not necessarily be possible to place every pair of alternatives on a common scale such as utility. And if not, then traditional maximizing approaches – which after all are mathematical functions that depend on some common scale of measurement – will fail to yield an account of rational choice.

This is not the place to survey all the different ways of understanding value incommensurability or of incorporating limited or widespread incommensurability into the account of rational choice. For now, it is enough to observe two quite distinct responses to value incommensurability. One perspective can be found in Michael Weber's paper (Chapter 5). In it, he argues that we can understand our lives from two distinct and incommensurable temporal perspectives: that of the moment and that of a whole life. In some instances, an alternative might be best from one perspective and suboptimal from the other. Weber contends that the incommensurability of the perspectives might yield a *rational permission* to satisfice with respect to one of them. This permission is supposed to parallel an agent-centered moral permission in moral theory to do other than what would maximize agent-neutral well-being. On some consequentialist accounts it is moral, for example, to save the life of my child in preference to saving two strangers, and we might account for this discrepancy from the model of maximizing agent-neutral value in terms of an agent-centered permission that conditions the demands of maximizing. Similarly in the theory of individual rational choice, the whole-life and momentary perspectives might condition each other, for instance with respect to career ambitions. One might, from the whole-life perspective, aspire to career greatness and yet at any particular moment be unwilling to sacrifice leisure, family, or other ends to do what greatness requires. In this instance the values of the momentary perspective condition those of the whole-life perspective, yielding a rational permission

to seek less than what is optimal. Incommensurability thus underwrites a form of rational satisficing, on Weber's account.

Henry Richardson (Chapter 6) also argues from a standpoint that recognizes value incommensurability, yet he reaches a conclusion quite different from Weber's. The heart of Richardson's argument is the observation that a satisficing conception of practical rationality places incompatible demands on the theory. On the one hand, the notion of "tradeoffs" implicit in the idea of settling for what is "good enough" suggests that satisficing will be a rational strategy only when applied to relatively local pursuits, such as buying a car or choosing one's clothes. One cannot satisfice with respect to one's global goal of pursuing the good, because the global context would not provide any constraints such that an alternative would be good enough. In the global context, one chooses the best option, though locally satisficing can be rational. That said, the theorist of satisficing must provide some metric along which to assess some alternatives as good enough and others as not that good. Ordinarily, the metric is utility or preference satisfaction, and these notions are quite global. Richardson thus identifies a tension within the theory of satisficing: Although applicable only to local ends, its metric is a global one. This global metric, moreover, runs afoul of value incommensurability of the sort we confront every day, according to Richardson, who has his own non-optimizing, non-satisficing theory of rationality. In this case, it seems, incommensurability is a reason to reject a satisficing theory.

No doubt it is too quick to oppose Weber and Richardson in this fashion. Weber addresses incommensurable perspectives, and Richardson treats incommensurable values. And yet their contributions can speak to each other: The value incommensurability that drives Richardson's account might be accommodated differently by Weber's diverse temporal perspectives. The latter's account of choice in terms of temporal perspectives might be built into different values by Richardson. Theorists of rational choice will benefit from these reflections on incommensurability and how to handle it.

Notes

1. H. A. Simon, "A Behavioral Model of Rational Choice," *Quarterly Journal of Economics* 69 (1955): 99–118.
2. Alternatively, one might adopt a three-valued function, $(1, 0, -1)$, corresponding to "win, draw, or lose." This nuance of Simon's discussion need not concern us.

3. For an excellent and straightforward introduction to expected utility theory, see Michael D. Resnik, *Choices*. Minneapolis: University of Minnesota Press, 1987.
4. Michael Slote, *Beyond Optimizing: A Study of Rational Choice*. Cambridge, Mass.: Harvard University Press, 1989.
5. Michael Slote, "Satisficing Consequentialism," *Proceedings of the Aristotelian Society* 58 supp. (1984): 139–163.
6. J. S. Mill, *Utilitarianism, Second Edition*. Indianapolis: Hackett, 2001, p. 7.
7. Notice that this argument does not depend on a so-called "revealed preference" account, according to which we reveal our actual preferences through our choices. Rather, it explains the initial expression of preference as perhaps only partial, and thus one that incompletely or inadequately characterizes the alternatives under consideration.
8. I develop this argument in Michael Byron, "Satisficing and Optimality," *Ethics* 109 (1998): 67–93.