personalities or training opportunities affect team performance. This brings to mind Wittgenstein's remark about psychology combining empirical methods and conceptual confusion. But even here there are less obvious results such as that the contribution of a high ability team member appears to be greater when other team members also have high ability.

In conclusion, the breadth of material covered and overall quality make this a very worthwhile collection for anyone interested in teamwork. The downside is that it provides a case study of the fragmentation of the social sciences. Whilst this is at times frustrating, the picture is doubtless accurate. Greater integration is surely furthered by efforts such as this to bring related fields together, though, and, as I hope to have indicated, there is much material here capable of cross-fertilizing different disciplines.

Nicholas Bardsley
National Centre for Research Methods, University of Southampton

REFERENCE

Bacharach, Michael. 2006. *Beyond Individual Choice*, ed. Nathalie Gold and Robert Sugden. Princeton University Press.

doi:10.1017/S0266267107001393

Satisficing and Maximizing: Moral Theorists on Practical Reason, edited by Michael Byron. Cambridge University Press, 2004, 245 pages.

Is only the best action good enough to choose? Maximizers would answer 'yes', and tell us to always choose the best action, whereas satisficers would answer 'no' and tell us to choose an action that is merely good enough. This way of putting things is still very rough, since it is not clear what 'good enough' means. It is one of the virtues of this book that it shows how ambiguous the notion is. In fact, many of the contributors of this book argue that on a natural understanding of 'good enough' even maximizers can happily accept that it is sometimes permissible to do what is good enough. Another virtue of this book is that it shows that the notion of doing what is good enough is not just relevant to rational choice narrowly conceived. In the hands of the contributors, the notion of 'good enough' is used to illuminate the virtue of moderation, the notions of supererogation, 'demandingness' and incommensurability, the relation between well-being and prudential choice, and the distinction between deontology and consequentialism.

In this short review, I will not be able to discuss all these issues. I will mainly focus on what I take to be the central question: Can satisficing be seen as a plausible alternative to maximizing? This means that I will

say very little about Slote's and Hurka's interesting contributions because they both assume an affirmative answer to this question and then go on to discuss the implications this has for the relation between satisficing and the putative goods of moderation and perfection.

The term 'satisficing' was coined by the economist Herbert Simon. As Byron points out in his informative introduction to the book, what motivated Simon to introduce this notion was that the standard maximizing form of decision theory seemed overly demanding. It tells us that an action is rational just in case it maximizes expected utility. But in order to decide which act maximizes expected utility I need to know my options, and, for each option, the precise utilities and probabilities of its possible outcomes. It seems unrealistic to assume that we have this knowledge. Simon suggested instead that we should judge the possible outcomes of an action as 'satisfactory' or 'unsatisfactory', and say that an action is rational just in case it guarantees a satisfactory outcome (i.e., just in case all the possible outcomes of the option will be satisfactory). Simon also suggested a 'stopping rule' that there is no need to contemplate further options once we have found one that will guarantee a satisfactory outcome.

This seems to take care of the 'problem of demandingness', but it is not so clear that this was a genuine problem to start with. For, as Dreier points out in his contribution to the volume, this assumes a very crude understanding of decision theory. It assumes that standard decision theory is in the business of telling us *how to deliberate* when we make decisions. But according to another interpretation, decision theory only tells us that rational choices can be *represented* by a utility function in such a way that the rational action turns out to be one that maximizes expected utility. More exactly, it can be proved that if the agent's preferences over outcomes (and lotteries of outcomes) satisfy certain axioms, then they can be represented by a utility function that makes his choices utility-maximizing. On this understanding of the theory, no advice is given to the agent on how to deliberate about what to do.

It is important to stress that the notion of utility employed here is just a numerical representation of preferences. Hence, decision theory does not require us to maximize some quantity, such as pleasure, well-being, or money. On this theory, there is simply no quantity called 'utility' that a person can identify and seek to maximize. Therefore, when in Slote's famous example you refrain from having a tasty snack that you believe would be enjoyable you may be perfectly rational according to decision theory, for it does not require you to maximize enjoyment.

Even though it is difficult to see satisficing as a genuine alternative to traditional decision theory, there are other ways to spell out this notion. In a trivial sense, we all are satisficers. No one would claim that just because the presence of a certain feature makes things better, we must opt for the maximum amount of this feature. The best amount need not be the

maximum amount. For instance, to take a trivial example, just because the presence of salt makes a dish better, it does not follow that we should bombard the dish with salt.

Even if more of a certain quantity is always better, it does not follow that we should maximize the amount of this quantity. There might be more than one evaluatively relevant quantity to consider and we need to strike a balance between two or more maximands. Schmidt, Narveson, Cowen, and Byron argue that in this case satisficing can be seen as a *cost-saving strategy*. For example, when you are looking for a new toothpaste to buy it would be silly to seek the perfect toothpaste even if you have enough time to find one. A satisfactory toothpaste is good enough given the other ends you want to pursue. Seeking the perfect toothpaste would mean leaving very little time for other more important ends. Similarly, turning to a more serious case, even if you care a lot about your own personal well-being, you can opt for a satisfactory level rather than the optimal because you also want to be able to promote the well-being of other people.

As Schmidtz and Narveson point out, satisficing can also be seen as *good decision method*. Searching for the perfect partner is not the best way to succeed in finding one, as Narveson reminds us. If you are constantly assessing your prospective partners' qualities, you will scare them away. It would be better to take an interest in the relationship itself because then you will have a greater chance of achieving your original end.

Common to these approaches to satisficing is that they can be easily adopted by maximizers. Is there a plausible form of satisficing that would deny that the right thing to do is always the best? To make this question interesting we need to assume that there is a best action to consider. Obviously, if there were no best action it would be foolish to demand that we do the best action. The more challenging question is whether it can be right to choose an action when you know that there is another option that is better.

As the contributions to this volume make clear, the answer to this question hinges crucially on what 'better' means. If it is equated with 'more reason to do', the answer seems obvious. I fully agree with Byron, Narveson, Richardson, and Schmidtz that it is impermissible to choose an action when you know that there is another action that you have overall more reason to do.

One way to evade this problem is to relativize reason to different perspectives. This is what Dreier does in his attempt to provide a coherent notion of moral supererogation. (He denies that this strategy works for Slote's notion of 'rational supererogation'.) Supererogatory actions seem paradoxical since it is morally permissible to refrain from doing a supererogatory action even though the balance of moral reason speaks in favour of doing it. Dreier's solution to this paradox is to deny that there is an all-things-considered moral perspective from which we can judge

the supererogatory action. What we have instead are two different moral perspectives that cannot be amalgamated. From the more demanding perspective of beneficence, you have most reason to do the supererogatory action and it would be wrong not to do it, whereas from the less demanding perspective of justice, you do not have any reason to perform this action and you are thus permitted to refrain from doing it.

This simple solution is not cost-free, however. Dreier has to accept that refraining to do the supererogatory action is wrong *from one moral perspective*. But the common intuition, I think, is rather that a failure to do this action is not a moral failure at all.

It is important to note that by relativizing permissibility and value in this sense we can accept a qualified version of one of the satisficer's commitments: it can be permissible, *from one perspective*, to do something even when there is another option that is better *from another perspective*. But it is doubtful whether this is enough to qualify as true satisficing. To be a true satisficer it does not seem enough to accept that what is permissible from one perspective might be suboptimal from another.

One obvious reason why this is not enough is that the perspectives might belong to *different agents* as in Swanton's favoured version of virtue ethics. She argues that the virtue of perfection need not be overly demanding since this virtue will require different things for different agents. All agents are required to do what is best, but what is permissible for an ordinary agent need not be what is best for a supremely virtuous agent. Swanton claims that this is a case where ordinary agents are permitted to satisfice in 'a weaker sense'. But this is seriously misleading since she admits that we are always required to do the best. Indeed, as Hurka convincingly argues, maximizing fits especially well with perfectionist goods. No one would suggest that the motto of the Olympics should be 'Reasonably fast, reasonably high, reasonably strong'.

Even when permissibility and value are relativized to different perspectives belonging to the same agent, we need not have a form of satisficing. A radical deontologist who claims that an action is permissible just in case it does not violate deontological constraints would of course accept that what is permissible from the perspective of deontological constraints need not be what is best from the perspective of agent neutral value. But, as van Roojens convincingly argues, this does not show that radical deontologists are satisficers. They simply deny that the agent-neutral values of options have any bearing on what is permissible. To be a true satisficer, it is also important that one accepts that an action may be permissible in part *because* it is good enough.

Weber in his contribution argues for this more substantive form of satisficing. He argues that there are cases where it is prudentially permissible to do what is not best for the agent but it is permissible in

part because it is good enough for the agent. Like Dreier, Weber makes sense of satisficing by adopting a relativization strategy, but, unlike Dreier, he only relativizes the values, not permissibility. He argues that there are two different perspectives from which we can assess the well-being of a person, a narrower and a broader. From the narrower perspective, we assess the well-being of a person in terms of her momentary well-being. From this perspective, even minor pains matter a lot. From the broader perspective, we assess the well-being of the life *as a whole* and from this perspective, minor pains have little or no significance. What matters from this perspective is the overall pattern of the life, for instance, whether early hardships are linked to later success. Weber denies that these two perspectives can be amalgamated into a more inclusive perspective from which we can make all-things-considered judgements about the well-being of a life. To decide what we have most prudential reason to do we need to consider both perspectives. In particular, we are prudentially permitted to do what is best in terms of momentary well-being given that it is still good enough from the perspective of one's life as a whole.

I welcome the distinction between a life's momentary well-being and a life's well-being as a whole, but I very much doubt that Weber is right in denying the possibility of aggregating these different sources of well-being. It is true that, as Weber points out, it would be a mistake to focus exclusively on the broader perspective when assessing a life since that would be to ignore completely the momentary, day-to-day, ups and downs. But there seems to be third, even broader, perspective from which we judge the value of a life. From this perspective, we judge the value of a life *on the whole*, and not as a whole. In order to assess the value of a life, on the whole we need to consider both a life's momentary well-being and the well-being it has as a whole. A life's value on the whole is some function (not necessarily additive) of all the good and bad patches in the life and the degree of unity the life exhibits. By failing to distinguish between value as a whole and value on the whole, Weber stacks the cards against this third option. It does not seem plausible to say that no lives can be compared in terms of the values they have on the whole. All momentary well-being being equal, a more coherent and unified life seems more valuable on the whole than a less coherent and unified one.

What is common to all these defences of satisficing is that the action that is deemed permissible is not suboptimal in terms of all normatively relevant aspects or from all normatively relevant perspectives. But perhaps it is not especially surprising that the right and the best can come apart, if you work with a notion of betterness that is not intended to capture everything that is relevant for your choice. It is true, as Richardson argues, that in these cases it could still be sensible to choose an option that is good enough in the sense of satisfying each of the incommensurable constraints we impose on the choice. But if all the options can be ordered by

a transitive and complete betterness relation that incorporates all relevant aspects or perspectives, there does not seem to be a plausible alternative to maximizing.

Krister Bykvist
Jesus College, Oxford

doi:10.1017/S026626710700140X

Economic Theory and Cognitive Science, by Don Ross. MIT Press, 2005, 384 pages.

Don Ross' *Economic Theory and Cognitive Science* is a challenging, well thought out book that exhibits considerable understanding of economics, philosophy, and cognitive science, and deserves to be taken seriously. Its premise that economics must not only address cognitive science but change in response to it seems entirely correct, though if current experience is any indication most economists will realize this only long after the die of change is cast. Of the many things worth discussing in the book, I will not address Ross' critique of eliminative materialism and intentional-stance functionalism alternative (cf. Nagel 1986; Searle 1997), his reading of the history of economics regarding Robbins and Samuelson, his view that reality boils down to a fundamental unity of one underlying kind of stuff (cf. Dupré 2001), his separateness of economics thesis (cf. Hausman 1992), nor his radical scientific realism and rejection of commonsense ontology (cf. Mäki 1992). I will address what I take to be the pivotal focus of the book, namely the ontological thesis that human individuals or selves are not agents but their subpersonal aspects are. I begin with a summary of Ross' relevant arguments, and then move to their evaluation.

RESCUING NEOCLASSICISM?

Ross seeks to unite the 'core insights of neoclassical economics with evolutionary cognitive and behavioral science' in a way that abandons both 'our conventional, "folk" schema for sorting intentional, behavioral and social reality' (19), and also the traditional assumption that human individuals or selves are agents. His approach is the opposite of that recommended by many other advocates of behavioral and evolutionary economics (e.g. Bowles 2003), who see the new programs as essentially anti-neoclassical. Ross rejects their position as relying on 'hyperempiricist methodological principles' and as a misguided attempt to transform economics into 'a branch of applied social psychology' (28), violating the purported status of economics as a separate science. Rather he argues that 'the core neoclassical commitment to economics as the systematic science